Avaya

Installation and Configuration Guide

AVAYA P333R-LB

STACKABLE SWITCH

SOFTWARE VERSION 4.0

April 2003

AVAYA

Table of Contents

Table of Contents i
Safety Information ix
FCC Notice ix
Conventions Used in the Documentation ix
CLI Conventions ix
Notes, Cautions and Warnings x

Section 1 Overview of the P330

Chapter 1 Avaya P333R-LB Overview 1
Introduction 1
Avaya P330 Family Features 1
Avaya P330 Network Management 2
Avaya P330 Device Manager (Embedded Web) 2
Avaya P330 Command Line Interface (CLI) 2
Avaya Multi-Service Network Manager™ (MSNM) 2
Port Mirroring 2
SMON 3
Fans, Power Supply and BUPS Monitoring 3

Chapter 2 Standards and Compatibility 5
Avaya P330 Standards Supported 5
IEEE 5
IETF - Layer 2 5
IETF - Layer 3 5
IETF - Network Monitoring 6

Chapter 3 Specifications 7
Avaya P333R-LB Switch 7
Physical 7
Power Requirements – AC 7
Power Requirements – DC 7
Environmental 7
Safety 8
Safety - AC Version 8
Safety - DC Version 8
Agency Approvals 8
EMC Emissions 8
Immunity 8

Other 9
Interfaces 9
Basic MTBF 9
Stacking Module 9
Expansion Modules 9
Gigabit Ethernet Expansion Modules 9
Laser Safety 9
Laser Classification 10
Usage Restriction 10
Laser Data 10
Fast Ethernet Fiber Expansion Module 11
Ethernet/Fast Ethernet Expansion Module 11
GBIC Expansion Module 12
Safety Information 12
Usage Restriction 12
Avaya Approved GBIC Transceivers 13
Specifications 13
Agency Approval 14
MTBF 14
X330GT2 Gigabit Ethernet Expansion Module 14
ATM Expansion Modules 14
Safety Information 15
WAN Expansion Modules 15

Section 2 Installing the P330

Chapter 4 Installation 19
Required Tools 19
Site Preparation 19
Rack Mounting (Optional) 21
Stacking Switches (Optional) 22
Installing the X330STK Stacking Sub-module in the P330 22
Connecting Stacking Sub-modules 22
Installing Expansion Sub-modules 25
Installing the Expansion Sub-module into the Avaya P330 25
Making Connections to Network Equipment 26
Prerequisites 26
Port Types 26

Chapter 5 Powering Up the Avaya P330 29
Powering On – Avaya P330 Module AC 29
Powering On – Avaya P330 Module DC 29
Post-Installation 30

Chapter 6 Avaya P333R-LB Front and Back Panels 31

Avaya P333R-LB Front Panel 31
BUPS Input Connector 34

Chapter 7 Establishing Switch Access 35
Establishing a Serial Connection 35
Configuring the Terminal Serial Port Parameters 35
Connecting a Terminal to the Avaya P330 Serial port 35
P330 Sessions 36
Assigning P330's IP Stack Address 36
Establishing a Telnet Connection 37
Establishing a Modem (PPP) Connection with the P330 38
Overview 38
Connecting a Modem to the Console Port 38

Chapter 8 User Authentication 39
Introduction 39
CLI Architecture 39
Security Levels 39
Entering the Supervisor Level 40
Defining new local users 40
Exiting the Supervisor Level 41
Entering the CLI 41
RADIUS 41
Introduction to RADIUS 41
Radius Commands 43
Allowed Managers 45
Allowed Manager CLI Commands 45

Section 3 Configuration of the P330

Chapter 9 Basic Switch Configuration 49
Introduction 49
System Parameter Configuration 50
Identifying the system 50
Operating parameters 50
Network Time Acquiring Protocols Parameter Configuration 51

Chapter 10 Default Settings of the P330 53
Configuring the Switch 53
Avaya P330 Default Settings 53

Chapter 11 Avaya P330 Layer 2 Features 57
Overview 57
Ethernet 57
Fast Ethernet 58

Gigabit Ethernet 58
Configuring Ethernet Parameters 58
Auto-negotiation 58
Full-Duplex/Half-Duplex 58
Speed 58
Flow Control 59
Priority 59
MAC Address 59
CAM Table 60
Ethernet Configuration CLI Commands 60
Ethernet Implementation in the Avaya P333R-LB 61
VLAN Configuration 62
VLAN Overview 62
VLAN Tagging 63
Multi VLAN Binding 63
Ingress VLAN Security 64
VLAN CLI Commands 64
VLAN Implementation in the Avaya P333R-LB 66
Port Based Network Access Control (PBNAC) 67
How "Port Based" Authentication Works 67
PBNAC Implementation in the P330 Family 67
Configuring the P330 for PBNAC 68
PBNAC CLI Commands 68
Spanning Tree Protocol 71
Overview 71
Spanning Tree Protocol 71
Spanning Tree per Port 71
Rapid Spanning Tree Protocol (RSTP) 72
About the 802.1w Standard 72
Port Roles 72
Spanning Tree Implementation in the P330 Family 73
Spanning Tree Protocol CLI Commands 74
MAC Security 76
MAC Security Implementation in P330 76
MAC Security CLI Commands 76
LAG 78
LAG Overview 78
LAG CLI Commands 78
LAG Implementation in the Avaya P330 Family of Products 79
Port Redundancy 80
Port Redundancy Operation 80
Intermodule Port Redundancy 81
Port Redundancy CLI Commands 81
IP Multicast Filtering 83

Overview 83
IP Multicast CLI Commands 84
IP Multicast Implementation in the Avaya P333R-LB 84
Stack Health 85
Overview 85
Implementation of Stack Health in the P330 Family 85
Stack Health CLI Commands 85
Port Classification 86
Overview 86
Port Classification CLI Commands 86
Stack Redundancy 87

Chapter 12 Avaya P330 Layer 3 Features 89
Introduction 89
What is Routing? 89
Routing Configuration 91
Forwarding 91
Multinetting (Multiple Subnets per VLAN) 91
IP Configuration 92
IP Configuration CLI Commands 92
Assigning Initial Router Parameters 93
RIP (Routing Interchange Protocol) Configuration 95
RIP Overview 95
RIP2 96
RIP CLI Commands 96
OSPF (Open Shortest Path First) Configuration 98
OSPF Overview 98
OSPF CLI Commands 99
Static Routing Configuration 100
Static Routing Overview 100
Static Routing Configuration CLI Commands 100
Route Preferences 101
Route Redistribution 102
Route Redistribution Commands 102
ARP (Address Resolution Protocol) Table Configuration 103
ARP Overview 103
The ARP Table 104
ARP CLI Commands 104
BOOTP/DHCP (Dynamic Host Configuration Protocol) Relay Configuration 105
BOOTP/DHCP Overview 105
BOOTP 105
DHCP 105
DHCP/BOOTP Relay 105
BOOTP/DHCP CLI Commands 106

NetBIOS Re-broadcast Configuration 107
NetBIOS Overview 107
NetBIOS Re-broadcast Configuration CLI Commands 107
..... 107
Policy Configuration 108
Policy Configuration Overview 108
Enforcement 109
Default List Behavior 109
Policy and Load Balancing 109
P333R-LB Policy Properties 109
Policy Configuration CLI Commands 109
Policy Configuration Example 111
Policy Configuration Example 111
IP Fragmentation and Reassembly 112
IP Fragmentation and Reassembly Overview 112
IP Fragmentation/Reassembly CLI Commands 112
Layer 3 Configuration File 113

Chapter 13 Layer 3 Redundancy 115
VRRP 115
VRRP Commands 116
Configuration Example 118
SRRP 121
SRRP Commands 121
Additional Redundancy Schemes 122
Real Server Group Backup 122
RSG Backup Configuration Example 122
Real Server Backup 125
Real Server Backup Configuration Example 125

Chapter 14 Load Balancing in the P333R-LB 1
Firewall Load Balancing 1
Benefits 1
How It Works 1
Transparent Routing Firewall Load Balancing 2
Implementation 2
Configuring Firewall Load Balancing in the P333R-LB 3
Demilitarized Zone (DMZ) Configuration Example 9
Load Balancing Metrics for Transparent Routing FWLB 13
Hash 13
MinMiss Hash 14
Selecting a Load Balancing Metric 15
Weighted Firewalls 15
Health Check 15
Persistency 16

Non-Transparent Routing Firewall Load Balancing 16
Implementation 16
NAT Configuration Example 17
Static NAT Configuration Example 18
Load Balancing Metrics for Non-Transparent Routing FWLB 19
Health Check 20
Persistency 20
Bridging Firewall Load Balancing 20
Implementation 20
Configuration Example 21
Load Balancing Metrics for Bridging FWLB 24
Health Check 24
Persistency 24
Firewall Load Balancing Management Security 25
Server Load Balancing..... 26
Benefits 26
How it Works 26
Half NAT Load Balancing 26
Configuring Server Load Balancing in the P333R-LB 28
Half NAT Based Configuration 28
Full NAT Load Balancing 30
Full NAT Based Configuration 32
Direct Server Return (DSR) (Triangulation) Redirection 34
DSR (Triangulation) Configuration Example 35
DNS Configuration 37
Combining Two Cache Applications on the Same Real Caches 37
Load Balancing Metrics 37
Round Robin 38
Hash 38
MinMiss Hash 39
Weighted Real Servers 40
Health Check 40
Client Persistency 41
Selecting a Load Balancing Metric and Persistency Option 42
Port Re-mapping 42
Application Redirection..... 44
Benefits 44
How It Works 44
Configuration Examples 46
Application Redirection 46
Combining Two Cache Applications on the Same Real Caches 49
Load Balancing Metrics 52

Round Robin 52
Hash 52
MinMiss Hash 53
Weighted Real Servers 54
Health Check 54
Client Persistency 56
Policy-Based Routing (Source-Based Routing) 57
Selecting a Load Balancing Metric and Persistency Option 57

Chapter 15 Embedded Web Manager 59
Overview 59
System Requirements 59
Running the Embedded Web Manager 60
Installing the Java Plug-in 63
Installing from the Avaya P330 Documentation and Utilities CD 63
Install from the Avaya Site 63
Install from your Local Web Site 63
Installing the On-Line Help and Java Plug-In on your Web Site 64

Section 2 Troubleshooting and Maintaining the P330

Chapter 16 Troubleshooting the Installation 67
Troubleshooting the Installation 67

Chapter 17 Maintenance 69
Introduction 69
Adding/Replacing an Expansion Sub-module 69
Adding an Expansion Sub-module to Avaya P330 69
Replacing an Existing Expansion Sub-module 69
Saving Configuration Definitions 70
Without Saving Configuration Definitions 70
Replacing the Stacking Sub-module 70

Chapter 18 Updating the Software 71
Software Download 71
Obtain Software Online 71
Downloading Software 71
Download New Version without Overwriting Existing Version 72

How to Contact Us 73
In the United States 73
In the AP (Asia Pacific) Region 74
In the CALA (Caribbean and Latin America) Region 75

Before you Install the P333R-LB

Safety Information

Caution: The Avaya P330 switch and modules contain components sensitive to electrostatic discharge. Do not touch the circuit boards unless instructed to do so.

Caution: Do not leave any slots open. Cover empty slots using the blanking plates supplied.

Warning: The fans are on whenever the power is on in the chassis.

FCC Notice

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Changes or modifications to this equipment not expressly approved by Avaya Inc. could void the user's authority to operate the equipment.

Conventions Used in the Documentation

Documentation for this product uses the following conventions to convey instructions and information:

CLI Conventions

- Mandatory keywords are in the **`computer bold`** font.

- Information displayed on screen is displayed in `computer` font.
- Variables that you supply are in pointed brackets <>.
- Optional keywords are in square brackets [].
- Alternative but mandatory keywords are grouped in braces {} and separated by a vertical bar |.
- Lists of parameters from which you should choose are enclosed in square brackets [ ] and separated by a vertical bar |.
- If you enter an alphanumeric string of two words or more, enclose the string in inverted "commas".

Notes, Cautions and Warnings

Note: Notes contain helpful information or hints or reference to material in other documentation.

Caution: You should take care. You could do something that may damage equipment or result in loss of data.

Warning: This means danger. Failure to follow the instructions or warnings may result in bodily injury. You should ensure that you are qualified for this task and have read and understood *all* the instructions

© 2003 Avaya Inc. All rights reserved. All trademarks identified by the ® or TM are registered trademarks or trademarks, respectively, of Avaya Inc. All other trademarks are the property of their respective owners.

AVAYA P333R-LB

SECTION 1: OVERVIEW OF THE P330

Chapter 1

Avaya P333R-LB Overview

Introduction

The Avaya P330 family of stackable Ethernet workgroup switches includes a range of modules with 10/100/1000 Mbps ports, a Layer 3 capability, and ATM and WAN expansion modules. The Avaya switch has 24x10/100 Mbps ports and an Expansion Module slot. The optional expansion modules provide additional Ethernet, Fast Ethernet, and Gigabit Ethernet connectivity.

An Avaya P330 stack can contain up to 10 switches and up to 3 backup power supply units. The stacked switches are connected using the Avaya X330STK stacking Modules which plug into a slot in the back of the Avaya P330. They are connected using the X330SC or X330LC cable (if the stack is split between two racks). The Avaya X330RC cable connects the top and bottom switches in the stack and provides redundancy and hot-swappability in the same way that modules can be swapped in a modular switching chassis.

The Avaya P330 is fully compliant with IEEE standards for VLAN Tagging, Gigabit Ethernet, Spanning Tree and Flow Control. This full standards-compliance, combined with auto-negotiation for 10/100/1000 Mbps and half/full duplex facilitates the expansion of your network to match your company's growing needs.

Avaya P330 Family Features

- You can connect up to 10 Avaya P330 switches in a stack. Moreover, this stack can be either in one rack or split over several racks using the X330LC Long Cable, according to your requirements.
- Avaya X330STK - this stacking Module is used to connect Avaya P330 switches in a stack, via the Octaplane.
- Avaya P330 BUPS - this back-up power supply module supports up to four Avaya P330 switches.
- One RJ-45/RS232 front panel console connector for both terminal and modem sessions.
- Two fan units in every switch, with operation sensors.
- One virtual IP address for managing the whole stack, the P330 stack is managed as a single entity.
- Hot-swapping of one switch at a time - by activation of the redundant cable:
  - — Does not disrupt the operation of other Avaya P330 switches.
  - — Does not change stack configuration.
  - — Does not require network downtime.

- Connection through Telnet from the front panel ports of *any* switch, with:
  — multiple levels of password protection
  — login and inactivity timeouts

Avaya P330 Network Management

Comprehensive network management is a key component of today's networks. Therefore we have provided multiple ways of managing the Avaya P330 to suit your needs.

Avaya P330 Device Manager (Embedded Web)

The built-in Avaya P330 Device Manager (Embedded Web Manager) allows you to manage an Avaya P330 stack using a Web browser without purchasing additional software. This application works with the Microsoft® Internet Explorer and Netscape® Navigator web browsers and Sun Microsystems Java™ Plug-in.

Avaya P330 Command Line Interface (CLI)

The Avaya P330 CLI provides a terminal type configuration tool for local or remote configuration of Avaya P330 features and functions.

Avaya Multi-Service Network Manager™ (MSNM)

When you need extra control and monitoring or wish to manage other Avaya network equipment, then the Avaya Multi-Service Network Manager network management suite is the answer. This suite provides the ease-of-use and features necessary for optimal network utilization.

Avaya Multi-Service Network Manager is available for Windows® NT®/2000 and Solaris 8. It can also operate in Stand-Alone mode with Windows® NT®/2000. Finally, Avaya Multi-Service Network Manager can operate under HP OpenView for Windows® NT®/2000 and Solaris 8.

Port Mirroring

The P330 provides port mirroring for additional network monitoring functionality. You can filter the traffic and mirror either incoming traffic to the source port or both incoming and outgoing traffic. This allows you to monitor the network traffic you need.

Ports which are members in a Link Aggregation Group (LAG) cannot *also* be used as Port Mirroring Destination or Source ports.

SMON

The P330 supports Avaya's ground-breaking SMON Switched Network Monitoring, which the IETF has now adopted as a standard (RFC2613). SMON provides unprecedented top-down monitoring of switched network traffic at the following levels:

- Enterprise Monitoring
- Device Monitoring
- VLAN Monitoring
- Port-level Monitoring

This top-down approach gives you rapid troubleshooting and performance trending to keep the network running optimally.

Note: MSNM Licence is required to run SMON monitoring.

Note: You need to purchase one SMON License per P330 Stack

Fans, Power Supply and BUPS Monitoring

The P330 module has integrated sensors which provide advance warnings of fan failure, power supply failure or Backup Power Supply (BUPS) failure via management.

Chapter 2

Standards and Compatibility

Avaya P330 Standards Supported

The Avaya P330 complies with the following standards.

IEEE

- 802.3x Flow Control on all ports
- 802.1Q VLAN Tagging support on all ports
- 802.1p Priority Tagging compatible on all ports
- 802.1D Bridges and STA
- 802.1w Rapid Spanning Tree Protocol
- 802.1X Port Based Network Access Control
- 802.3z Gigabit Ethernet on expansion module

IETF - Layer 2

- MIB-II - RFC 1213
- Structure and identification of management information for TCP/IP-based Internet - RFC 1155
- Simple Network Management Protocol (SNMP) - RFC 1157
- PPP Internet Protocol Control Protocol (IPCP) - RFC 1332
- PPP Authentication Protocols (PAP & CHAP) - RFC 1334
- PPP - RFC 1661
- ATM Management - RFC 1695
- RMON - RFC 1757
- SMON - RFC 2613
- Bridge MIB Groups - RFC 2674 dot1dbase and dot1dStp fully implemented. Support for relevant MIB objects: dot1q (dot1qBase, dot1qVlanCurrent)
- The Interfaces Group MIB - RFC 2863
- Remote Authentication Dial In User Service (RADIUS) - RFC 2865

IETF - Layer 3

- Internet Protocol - RFC 791
- Internet Control Message Protocol - RFC 792
- Ethernet Address Resolution Protocol - RFC 826
- Standard for the transmission of IP datagrams over Ethernet - RFC 894
- Broadcasting Internet datagrams in the presence of subnets - RFC 922

- Internet Standard Subnetting Procedure - RFC 950
- Bootstrap Protocol - RCF 951
- Using ARP to implement transparent subnet gateways - RFC 1027
- Routing Information Protocol - RCF 1058
- Hosts Extensions for IP Multicasting - RFC 1112
- Requirements for Internet Hosts - Communications Layers - RFC 1122
- DHCP Options and BOOTP Vendor Extensions - RFC 1533
- Interoperation between DHCP and BOOTP - RFC 1534
- Dynamic Host Configuration Protocol - RFC 1541
- Clarifications and Extensions for the Bootstrap Protocol Information - RFC 1542
- OSPF Version 2 - RFC 1583
- RIP Version 2 Carrying Additional Information - RFC 1723
- RIP Version 2 MIB Extension - RFC 1724
- Requirements for IP Version 4 Routers - RFC 1812
- OSPF Version 2 Management Information Base - RFC 1850
- IP Forwarding Table MIB - RFC 2096
- Virtual Router Redundancy Protocol - RFC 2338

IETF - Network Monitoring

- RMON (RFC 1757) support for groups 1,2,3 and 9
  - — Statistics
  - — History
  - — Alarms
  - — Events
- SMON (RFC 2613) support for groups
  - — Data Source Capabilities
  - — Port Copy
  - — VLAN and Priority Statistics
- Bridge MIB Groups - RFC 2674
  - — dot1dbase and dot1dStp fully implemented.
  - — Support for relevant MIB objects: dot1q (dot1qBase, dot1qVlanCurrent)

Chapter 3

Specifications

Avaya P333R-LB Switch

Physical

Height	2U (88 mm, 3.5″)
Width	482.6 mm (19″)
Depth	450 mm (17.7″)
Weight	7.5 kg (16.5 lb)

Power Requirements — AC

Input voltage	100 to 240 VAC, 50/60 Hz
Power dissipation	150 W max
Input current	5.3 A

Power Requirements — DC

Input voltage	-36 to -72 VDC
Power dissipation	150 W max
Input current	5.1 A max

Environmental

Operating Temp.	-5 to 50°C (23 to 122°F)
Relative Humidity	5% to 95% non-condensing

Safety

- UL for US approved according to UL195O Std.
- C-UL(UL for Canada) approved according to C22.2 No.950 Std.
- CE for Europe approved according to EN 60950 Std.
- Laser components are Laser Class I approved:
  - EN-60825/IEC-825 for Europe
  - FDA CFR 1040 for USA

Safety - AC Version

- Overcurrent Protection: A readily accessible Listed safety-approved protective device with a 16A rating must be incorporated in series with building installation AC power wiring for the equipment under protection.

Safety - DC Version

- Restricted Access Area: This unit must be installed in Restricted Access Areas only.
- Installation Codes: This unit must be installed in accordance with the US National Electrical Code, Article 110 and the Canadian Electrical Code, Section 12.
- Conductor Ampacity: Per UL 1950, Annex NAE (NEC Article 645-5(a)), the branch-circuit conductors supply shall have the ampacity of not less than 125 percent of the total connected load. For input leads use at least 18 AWG copper conductors.
- Overcurrent Protection: Per UL 1950, Annex NAE (NEC Article 240-3), a readily accessible listed branch-circuit overcurrent protective device rated maximum 10A must be incorporated into the building wiring.

Agency Approvals

EMC Emissions

Approved according to:

- US - FCC Part 15 Subpart B, Class A
- EU - EN55022 Class A
- EU - EN61000-3-2
- Japan - VCCI-A

Immunity

Approved according to:

- EN55024
- EU - EN61000-3-3

Other

Approved according to:

- CLEI Code: According to Tecordia (Bellcore) KS-22022 Standard
- NEBS Level 3 (optional mounting brackets)

Interfaces

- .
- RS-232 for terminal setup via RJ45 connector on front panel.

Basic MTBF

- hrs minimum

Stacking Module

Table A.1 Stacking Module

Name	Number of Ports
X330STK	2

Expansion Modules

Gigabit Ethernet Expansion Modules

Table A.2 Gigabit Ethernet Expansion Modules

Name	Number of Ports	Interface
X330S2	2	1000Base-SX
X330L2	2	1000Base-LX
X330S1	1	1000Base-SX
X330L1	1	1000Base-LX

Laser Safety

The Avaya X330S1/S2 multi-mode transceivers and the Avaya X330L1/X330L2

single mode transceivers are Class 1 laser products.

They comply with IEC 825-1 and Food and Drug Administration (FDA) 21 CFR 1040.10 and 1040.11.

The transceivers must be operated under recommended operating conditions.

Laser Classification

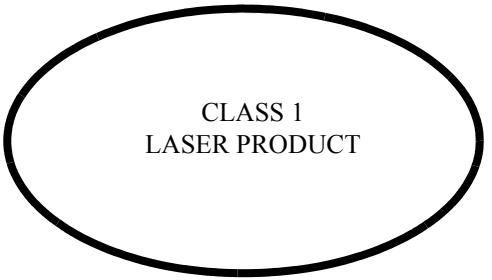

Note: Class 1 lasers are inherently safe under reasonably foreseeable conditions of operation.

Caution: The use of optical instruments with this product will increase eye hazard.

Usage Restriction

The optical ports of the module must be terminated with an optical connector or a dust plug when not in use.

Laser Data

Avaya P330S1/2 Expansion Modules

Wavelength: 850 nm

Output power dissipation: Max. 0.63W

Transmit power: Min. -9 dbm, Max. -4 dbm

Receive power: Min. -17 dbm, Max. 0 dbm

Avaya P330L1/2 Expansion Modules

Wavelength: 1300 nm

Output power dissipation: Max. 0.68W

Transmit power (9 µm SMF): Min. -9.5 dbm, Max. -3 dbm

Transmit power (62.5 µm and 50 µm MMF): Min. -11.5 dbm, Max. -3 dbm
Receive power (9 µm SMF, 62.5 µm and 50 µm MMF): Min. -20 dbm, Max. -3 dbm

Fast Ethernet Fiber Expansion Module

Table A.3 Fiber Fast Ethernet Expansion Module

Name	Number of Ports	Interface
X330F2	2	100Base-FX

Ethernet/Fast Ethernet Expansion Module

Table A.4 Ethernet/Fast Ethernet Expansion Module

Name	Number of Ports	Interface
X330T16	16	10/100Base-T

GBIC Expansion Module

The Avaya X330G2 Expansion Module is the GBIC (1.25 Gbit/s Gigabit Ethernet) Expansion Module for the Avaya P330 family of stackable switches.

Note: In order to use this module the Avaya P330 switch must must have Embedded S/W Version 2.2 or higher.

The X330G2 can be used either as a Gigabit Ethernet link or as a high Bandwidth backplane for connecting switches. The introduction of the GBIC interface to the Avaya P330 family presents an added value over the existing Gigabit Ethernet expansion modules. You can insert any of the Avaya-authorized GBIC transceivers into the X330G2 Expansion Module socket. This provides you with a highly modular and customisable Gigabit Ethernet interface. The GBIC transceivers are hot-swappable.

Safety Information

The multimode and single-mode GBIC transceivers are Class 1 Laser products. They comply with EN 60825-1 and Food and Drug Administration (FDA) 21 CFR 1040.10 and 1040.11.

The GBIC transceivers must be operated under recommended operating conditions.

Laser Classification

Note: Class 1 lasers are inherently safe under reasonably foreseeable conditions of operation.

Caution: The use of optical instruments with this product will increase eye hazard.

Usage Restriction

When a GBIC transceiver is inserted into the X330G2 Expansion Module but is not in use, then the Tx and Rx ports should be protected with an optical connector or a dust plug.

Avaya Approved GBIC Transceivers

Caution: All Avaya approved GBICs are 5V. Do not insert a 3.3V GBIC.

Avaya supplies the following two GBIC transceivers for the Avaya P330 X330G2 Expansion Modules. You can order these directly from your local Avaya representative using the PEC or COM Codes:

Type	Description	PEC Code	COM Code
GBIC SX Transceiver	Multimode Fiber 1000BaseSx (550 m)	4705-122	108659228
GBIC LX Transceiver	Single-mode Fiber 1000BaseLx (10 km)	4705-121	108659210

In addition, Avaya has tested and approved a number of GBIC transceivers from other manufacturers for use with the Avaya X330G2 Expansion Module. An up-to-date list can be found in Avaya's website at the following address: www.avaya.com/support

Specifications

X330G2- LX GBIC Transceiver

A 9 mm or 10 mm single-mode fiber (SMF) cable may be connected to a 1000Base-LX GBIC port. The maximum length is 10 km (32,808 ft).

A 50 mm or 62.5 mm multimode (MMF) fiber cable may be connected to a 1000Base-LX GBIC port. The maximum length is 550 m (1,804 ft.) for 50 mm and 62.5 mm cable.

The LX transceiver has a Wavelength of 1300 nm, Transmission Rate of 1.25 Gbps and Input Power of 5V.

X330G2- SX GBIC Transceiver

A 50 µm or 62.5 µm multimode (MMF) fiber cable may be connected to a 1000Base-SX GBIC port. The maximum length is 500 m (1,640 ft.) for 50 µm cable and 220 m (722 ft.) for 62.5 µm cable.

The SX transceiver has a Wavelength of 850 nm, Transmission Rate of 1.25 Gbps and Input Power of 5V.

Agency Approval

The transceivers comply with:

- EMC Emission: US – FCC Part 15, Subpart B, Class A;
  Europe – EN55022 class A
- Immunity: EN50082-1
- Safety: UL for US UL 1950 Std., C-UL (UL for Canada) C22.2 No.950 Std., Food and Drug Administration (FDA) 21 CFR 1040.10 and 1040.11, and CE for Europe EN60950 Std. Complies with EN 60825-1.

MTBF

The Mean Time Between Failures (MTBF) for the X330G2 Expansion Sub-module is 594,639 hours.

X330GT2 Gigabit Ethernet Expansion Module

The X330GT2 Expansion Module provides two copper Gigabit Ethernet 1000Base-T ports.

Note: The X330GT2 module is only supported by Avaya P330 embedded software versions 2.4 and higher.

ATM Expansion Modules

There are two Avaya P330 ATM Expansion Modules:

- X330-OC12F1: 500m, Multimode fiber, can also be OC-3 reduced range
- X330-OC12S1: 15 km, Single-mode fiber, can also be OC-3

The ATM Modules can be installed in the following Avaya P330 Family switches:

- Avaya P333T Hardware Version C/S 1.3 and higher, with Embedded S/W 2.4 and higher.

Note: The ATM Expansion Module cannot be used in Avaya P333T hardware Versions lower than C/S 1.3.

- Avaya P334T Embedded S/W Ver. 2.4 and higher.
- Avaya P332MF Embedded S/W Ver. 3.0 and higher.
- Avaya P333R Embedded S/W Ver. 2.4 and higher.

Refer to the Avaya X330 ATM Access Module Installation Guide for installation procedures.

The multimode Avaya X330-OC12F1 and X330-OC3F1 (future) ATM Modules are Class 1 LED products. The single-mode X330-OC12S1 ATM Module is a Class 1

Laser product. They comply with EN 60825-1 and Food and Drug Administration (FDA) 21 CFR 1040.10 and 1040.11.

The Modules must be operated under recommended operating conditions.

Safety Information

Single-mode Module Laser Classification

Note: Class 1 lasers are inherently safe under reasonably foreseeable conditions of operation.

Caution: The use of optical instruments with this product will increase eye hazard.

Multi-Mode Module LED Warning

The following warnings apply to the X330 ATM Modules equipped with multi-mode fiber.

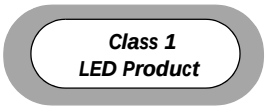

Warning: Class 1 LED Product. Do not view the LED through any magnifying device while it is powered on. Never look directly at the fiber Tx port and fiber cable ends when powered on.

WAN Expansion Modules

Avaya X330WAN is a series of WAN Edge Router expansion modules for the P330 Stackable Switching System . X330WAN enables you to connect your Avaya P330 switch to a WAN. X330WAN is part of Avaya's Converged Networks Solution that includes IP telephones, data switches and IP exchanges.

The X330WAN family includes the following modules:

- X330W-2DS1 access router module has 2 E1/T1 interfaces, a single 10/

100Base-T Fast Ethernet port, and a Console port.

- The X330W-2USP contains 2 USP (Universal Serial Ports), one 10/100Base-T Fast Ethernet port and one Console port.

An Avaya P330 stack can have X330WAN access router modules inserted in each of the switches in the stack with an expansion slot. A maximum stack configuration of 10 P334T switches using the X330WAN provides 490 Fast Ethernet 10/100 ports, and 20 E1/T1 or USP ports.

AVAYA P333R-LB

SECTION 2: INSTALLING THE P330

Chapter 4

Installation

This chapter describes the basic hardware Installation procedures for the Avaya P330.

Required Tools

Make sure you have the following tools at hand before undertaking the Installation procedures:

- Philips (cross-blade) screwdriver

Site Preparation

Avaya P330 can be mounted alone or in a stack in a standard 19-inch equipment rack in a wiring closet or equipment room. Up to 10 units can be stacked in this way. When deciding where to position the unit, ensure that:

- It is accessible and cables can be connected easily and according to the configuration rule.
- Cabling is away from sources of electrical noise such as radio transmitters, broadcast amplifiers, power lines and fluorescent lighting fixtures.
- Water or moisture cannot enter the case of the unit.
- There is a free flow of air around the unit and that the vents in the sides of the case are not blocked.

Note: Use Octaplane cables to interconnect with other switches.

- The environmental conditions match the requirements listed below:

Table 4.1 Environmental Prerequisites

Operating Temp.	-5 to 50°C (23 to 122°F)
Relative Humidity	5% to 95% non-condensing

- The power source matches the specifications listed below:

Table 4.2 Power Requirements — AC

Input voltage	100 to 240 VAC, 50/60 Hz

Power dissipation	150 W max
Input current	5.3 A

Table 4.3 Power Requirements — DC

Input voltage	-36 to -72 VDC
Power dissipation	150 W max
Input current	5.1 A max

Rack Mounting (Optional)

The Avaya P330 case fits in most standard 19-inch racks. Avaya P330 is 2U (88mm, 3.5″) high.

Place the Avaya P330 in the rack as follows:

1 Snap open the hinged ends of the front panel to reveal the fixing holes.

2 Insert the unit into the rack. Ensure that the four Avaya P330 screw holes are aligned with the rack hole positions as shown in Figure 4.1.

Figure 4.1 Avaya P330 Rack Mounting

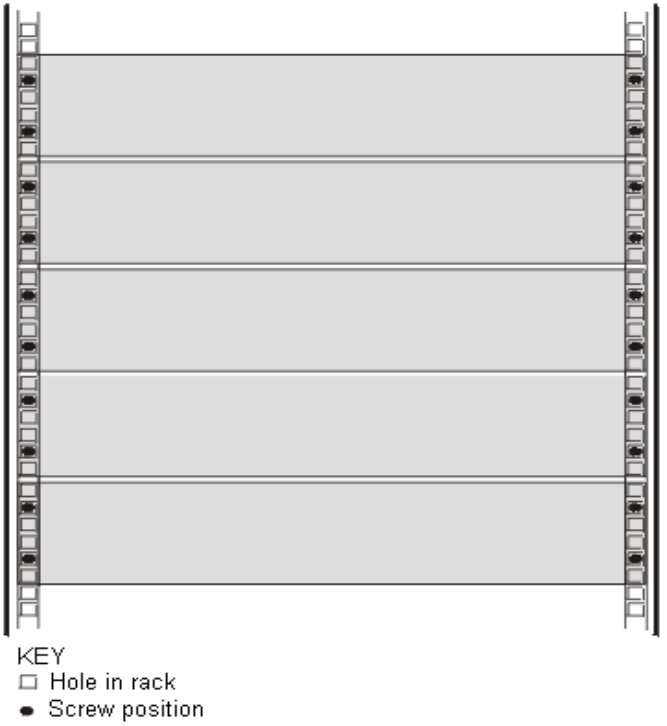

3 Secure the unit in the rack using the screws. Use two screws on each side. Do not overtighten the screws.

4 Snap closed the hinged ends of the front panel.

5 Ensure that ventilation holes are not obstructed.

Stacking Switches (Optional)

Avaya P330 is a stackable switching system. Stacking involves the mounting and connecting of stacking sub-modules in the P330 switch.

Installing the X330STK Stacking Sub-module in the P330

Caution: The stacking sub-modules contain components sensitive to electrostatic discharge. Do not touch the circuit board unless instructed to do so.

To install the stacking sub-module in the Avaya P330:

1 Remove the blanking plate from the back of the Avaya P330 switch.
2 Insert the stacking sub-module gently into the slot, ensuring that the metal base plate is aligned with the guide rails.
The metal plate of the X330STK (and *not* the PCB) fits onto the guide rails.
3 Press the sub-module in firmly until it is completely inserted into the Avaya P330.
4 Gently tighten the two screws on the side panel of the stacking sub-module by turning them.

Note: The Avaya P330 switch must not be operated with the back-slot open; the stacking sub-module should be covered with the supplied blanking plate if necessary.

Connecting Stacking Sub-modules

Before attempting to connect stacking sub-modules, verify that you have the required Octaplane cables.

Note: The two ends of the Octaplane cable terminate with different connectors. Each connector can only be connected to its matching port.

The following cables are used to connect stacked switches:

- Short Octaplane cable (X330SC) – ivory-colored, used to connect adjacent switches (Catalog No. CB0223) or switches separated by a BUPS unit.
- Long/Extra Long Octaplane cable (X330LC/X330L-LC) – ivory-colored, used to connect switches from two different physical stacks, or switches separated by a BUPS unit (Catalog No. CB0225/CB0270).
- Redundant/Long Redundant Octaplane cable (X330RC/X330L-RC) – black, used to connect the top and bottom switches of a stack (Catalog No. CB0222/CB0269).

These are the same cables that are used with all P330 family modules.

To connect stacked switches:

Note: When adding a module to an existing stack, first connect the stacking cables and then power up the module.

1 Plug the light grey connector of the Short Octaplane cable into the port marked "to upper unit" of the bottom Avaya P330 switch.
2 Plug dark grey connector of same Short Octaplane cable to the port marked "to lower unit" in the unit above. The connections are illustrated in Figure 4.3.
3 Repeat Steps 1 and 2 until you reach the top switch in the stack.
4 If you wish to implement stack redundancy, use the Redundant Cable to connect the port marked "to lower unit" on the bottom switch to the port marked "to upper unit" on the top switch of the stack.
5 Power up the added modules.

Caution: Do not cross-connect two Avaya P330 switches with two Octaplane (light-colored) cables. If you wish to cross-connect for redundancy, use one light-colored Octaplane cable and one black redundancy cable. Figure 4.2 shows an incorrect connection.

Note: You can build a stack of up to 10 Avaya P330 switches. If you do not wish to stack all the switches in a single rack, use long Octaplane cables to connect two physical stacks as shown in Figure 4.3.

Figure 4.2 Incorrect Stack Connection

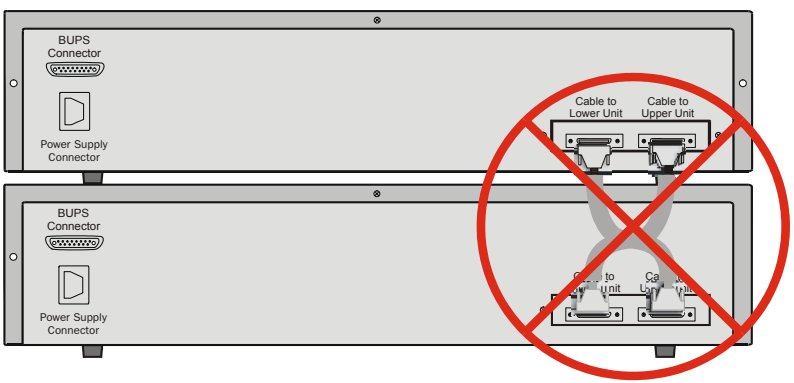

Figure 4.3 Avaya P330 Stack Connections

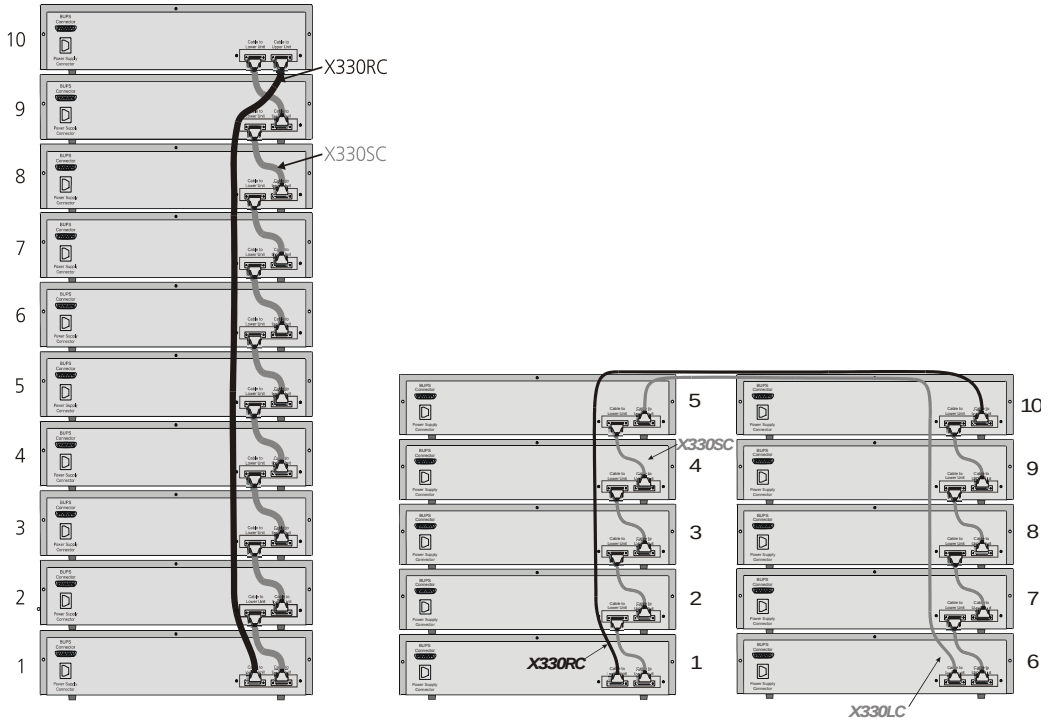

Installing Expansion Sub-modules

Caution: The expansion sub-modules contain components sensitive to electrostatic discharge. Do not touch the circuit board unless instructed to do so.

Installing the Expansion Sub-module into the Avaya P330

1 Remove the blanking plate or other sub-module (if installed).
2 Insert the sub-module gently into the slot, ensuring that the Printed Circuit Board (PCB) is aligned with the guide rails.
The PCB *not* the metal base plate fits into the guide rail.
3 Firmly press the sub-module until it is completely inserted into the Avaya P330.
4 Gently tighten the two screws on the front panel of the expansion sub-module by turning them.

Note: The Avaya P330 switch must not be operated with the expansion slot open; the expansion sub-module slot should be covered with the supplied blanking plate if necessary.

Making Connections to Network Equipment

This section describes the physical connections that you can make between the Avaya P330 switch and other network equipment.

Prerequisites

Make sure you have the following before attempting to connect network equipment to the P330 switch:

- a list of network equipment to be connected to the P330 switch, detailing the connector types on the various units
- all required cables (see below). Appropriate cables are available from your local supplier.

Port Types

Avaya P330 supports the following types of ports (according to the speed and standard they support):

- LAN — 10/100Base-T, 100Base-FX, 1000Base-T 1000Base-SX and 1000Base-LX
- WAN — by type:
  - — X330W-2DS1: E1/T1, 10/100Base-T
  - — X330W-2USP: USP (V.35), 10/100Base-T

Note: To interconnect Avaya P330 switches with twisted pairs, crossed cables are required.

- The maximum UTP cable length connected to a 10/100 Mbps port operating as 10Base-T, is 100 m (328 ft.).
- A UTP Category 5 cable must be connected to any 100Base-TX port, via an RJ45 connector. The maximum UTP cable length connected to a 10/100 Mbps port operating as 100Base-TX, is 100 m (328 ft.).
- A fiberoptic cable must be connected to any 100Base-FX port, via a SC connectors. The maximum fiber cable length connected to a 100Base-FX port is 412 m (1,352 ft) when operating in half duplex, and 2 km (6,562 ft) when operating in full duplex.
- A fiberoptic cable must be connected to 1000Base-SX or 1000Base-LX port, via SC connectors, according to the table below.

Table 4.4 Gigabit Ethernet Cabling

Gigabit Interface	Fiber Type	Diameter (µm)	Modal Bandwidth (MhzKm)	Maximum Distance (m)	Minimum Distance (m)	Wavelength (nm)
1000BASE-SX	MM	62.5	160	220	2	850
1000BASE-SX	MM	62.5	200	275	2	850
1000BASE-SX	MM	50	400	500	2	850
1000BASE-SX	MM	50	500	550	2	850
1000BASE-LX	MM	62.5	500	550	2	1310
1000BASE-LX	MM	50	400	550	2	1310
1000BASE-LX	SM	9	NA	10,000	2	1310

Powering Up the Avaya P330

This section describes the procedures for powering up the Avaya P330 unit.

Powering On – Avaya P330 Module AC

For the AC input version of the Avaya P330, insert the AC power cord into the power inlet in the back of the unit. The unit powers up.

If you are using a BUPS, insert a power cord from the BUPS into the BUPS connector in the back of the unit. The unit powers up even if no direct AC power is applied to the unit.

After power up or reset, the Avaya P330 performs a self test procedure.

applied to it.

Powering On – Avaya P330 Module DC

For the DC input version of the Avaya P330, connect the power cable to the switch at the input terminal block.

1 The terminals are marked "+", "-" and with the IEC 5019a Ground symbol.
2 The size of the three screws in the terminal block is M3.5.
3 The pitch between each screw is 9.5mm.

Connect the power cable to the DC power supply. After power up or reset, the Avaya P330 performs a self test procedure.

Warning: Before performing any of the following procedures, ensure that DC power is OFF.

Caution: This product is intended for installation in restricted access areas and is approved for use with 18 AWG copper conductors only. The installation must comply with all applicable codes.

Warning: The proper wiring sequence is ground to ground, positive to positive and negative to negative. Always connect the ground wire first and disconnect it last.

Post-Installation

The following indicate that you have performed the installation procedure correctly:

Table 5.1 Post-Installation Indications

Procedure	Indication	Troubleshooting Information
Powering the P330	All front panel LEDs illuminate briefly	Page 67
Creating Stacks	The LED next to the appropriate connection (“Cable to upper unit” or “Cable to lower unit”) is lit.	Page 67
Installing Expansion Modules	The LEDs on the Expansion Module flash briefly.	Page 67

If you do not receive the appropriate indication, please refer to "Troubleshooting the Installation".

Avaya P333R-LB Front and Back Panels

Avaya P333R-LB Front Panel

The Avaya front panel contains LEDs, controls, connectors and an expansion Module slot, as well as a console connector. The status LEDs and control buttons provide at-a-glance information.

The front panel LEDs consist of Port LEDs and Function LEDs. The Port LEDs display information for each port according to the illuminated function LED. The function is selected by pressing the left or right button until the desired parameter LED is illuminated.

Figure 6.1 shows the Avaya front panel. shows a detailed view of the LEDs (described in Table 6.1), pushbuttons, the Expansion Module slot, and the RJ-45 console connector at the bottom right.

Figure 6.1 P333R-LB Front Panel

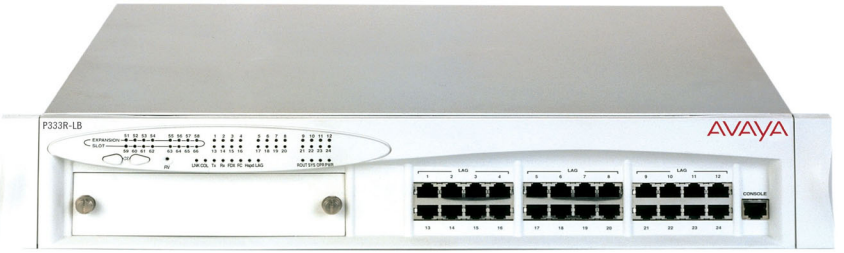

Figure 6.2 P333R-LB LEDs

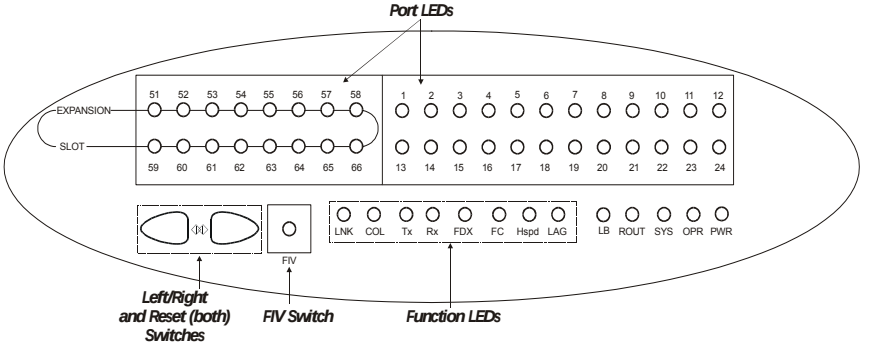
Port LEDs
EXPANSION
SLOT
51 52 53 54 55 56 57 58
59 60 61 62 63 64 65 66
1 2 3 4 5 6 7 8 9 10 11 12
13 14 15 16 17 18 19 20 21 22 23 24
FIV
LNK COL Tx Rx FDX FC Hspd LAG
LB ROUT SYS OPR PWR
Left/Right and Reset (both) Switches
FIV Switch
Function LEDs

Note: All LEDs are lit during a reset.

Table 6.1 Avaya LED Descriptions

LED Name	Description	LED Status
PWR	Power status	OFF – power is off
		ON – power is on
		Blink – using BUPS only
OPR	CPU operation	OFF – Module is booting
		ON – Normal operation
SYS	System Status	OFF – Module is a slave in a stack
		ON – Module is the Master of the stack and the Octaplane and Redundant cable are connected correctly. This LED will also light in Standalone mode.
		Blink – Box is the stack Master and the stack is in redundant mode.
The following Function LEDs apply to ports 1 to 66		
LNK	Port status	OFF – Port disabled
		ON – Port enabled and link OK
		Blink – Port enabled and the link is down
COL	Collision	OFF – No collision or FDX port
		ON – Collision occurred on line
Tx	Transmit to line	OFF – No transmit activity
		ON – Data transmitted on line from the module
Rx	Receive from line	OFF – No receive activity
		ON – Data received from the line into the module
FDX	Half/Full Duplex	OFF – Half duplex mode
		ON – Full duplex mode

Table 6.1 Avaya LED Descriptions

LED Name	Description	LED Status
FC	Flow Control	OFF – No Flow Control
		ON – Symmetric/Asymmetric Flow Control mode is *enabled* and port is in full duplex mode.
Hspd	High Speed	10/100 1000 OFF: 10 N/A ON: 100 1000
LAG	Link Aggregation Group (Trunking)	OFF – No LAG defined for this port
		ON – Port belongs to a LAG

Table 6.2 Avaya P330 <- -> Select buttons

Description	Function
Left/Right	Individual – select LED function (see table above).
Reset module	Press both right and left buttons together for approximately two seconds. All LEDs on module light up until buttons are released.
Reset stack	Press both right and left buttons together for 4 seconds. All LEDs on stack light up until buttons are released.
FIV	Not in use.

Note: The Port LEDs of the are numbered from . Expansion Module ports are numbered from 51. Port LED numbers 49-50 are reserved.

BUPS Input Connector

The BUPS input connector is a 5 VDC connector for use with the Avaya P330 BUPS unit only. A BUPS Input sticker appears directly to the right the BUPS input connector.

Figure 6.3 BUPS Input Connector Sticker

Chapter 7

Establishing Switch Access

This chapter describes various methods for accessing the Avaya P330 CLI, including:

- a terminal to the serial port on the switch
- P330 Sessions
- a workstation running a Telnet session connected via the network
- a remote terminal/workstation attached via a modem (PPP connection)

Establishing a Serial Connection

This section describes the procedure for establishing switch access between a terminal and the Avaya P330 switch over the serial port provided on the front panel of the P330 (RJ-45 connector labeled "Console").

Configuring the Terminal Serial Port Parameters

The serial port settings for using a terminal or terminal emulator are as follows:

- Baud Rate - 9600 bps
- Data Bits - 8 bits
- Parity - None
- Stop Bit - 1
- Flow Control - None
- Terminal Emulation - VT-100

Connecting a Terminal to the Avaya P330 Serial port

Perform the following steps to connect a terminal to the Avaya P330 Switch Console port for acessing the text-based CLI:

1 The P330 device is supplied with a console cable and a RJ-45-to-DB-9 adaptor. Use these items to connect the serial (COM) port on your PC/terminal to the Avaya P330 console port.

2 Ensure that the serial port settings on the terminal are 9600 baud, 8 bits, 1 stop bit and no parity.

3 When you are prompted for a Login Name, enter the default login. The default login is **`root`**.

4 When you are promoted for a password, enter the user level password **`root`**.

Now you can establish a connection to the Router or the Master switch (indicated when the SYS front panel LED is ON) using the Session commands and begin the configuration of Module, Stack, or Router parameters.

P330 Sessions

You can use sessions to switch between the CLI of P330 modules / other stack entities (for example, an X330 ATM or WAN entity plugged into a specific P330 switch or with the G700 Media Gateway Precessor) or to switch between Layer 2 and Layer 3 commands in the router module.

To switch between P330 modules use the command:
`session [<mod_num>] <mode>`.

The `<mod_num>` is the number of the module in the stack, counting from the bottom up.

The `<mode>` can be either **switch**, **router, wan, atm, mgp**.

Use **switch** mode to configure layer 2 commands.

Use **router** mode to configure routing commands.

Examples:

To configure router parameters in the module that you are currently logged into, type the following command:

`session router`.

To configure the switch parameters, on module 6, type the command:
`session 6 switch`.

Note: When you use the `session` command the security level stays the same.

Assigning P330's IP Stack Address

Note: All P330 switches are shipped with the same default IP address. You must change the IP address of the master P330 switch in a stack in order to guarantee that the stack has its own unique IP address in the network.

The network management station or a workstaion running Telnet session can establish communications with the stack once this address had been assigned and the stack has been inserted into the network. Use the CLI to assign the P330 stack an IP address and net mask.

To assign a P330 IP stack address:

1 Establish a serial connection by connecting a terminal to the Master P330 switch of the stack.
2 When prompted for a Login Name, enter the default name **`root`**
3 When you are prompted for a password, enter the password **`root.`** You are now in Supervisor Level.

4 At the prompt, type:
set interface inband <vlan> <ip_address> <netmask>
Replace <vlan>, <ip_address> and <netmask> with the VLAN, IP address and net mask of the stack.
5 Press Enter to save the IP address and net mask.
6 At the prompt, type **reset** and press Enter to reset the stack. After the Reset, log in again as described above.
7 At the prompt, type **set ip route** <dest> <gateway> and replace <dest> and <gateway> with the destination and gateway IP addresses.
8 Press Enter to save the destination and gateway IP addresses.

Establishing a Telnet Connection

Perform the following steps to establish a Telnet connection to the Avaya P330 for configuration of Stack or Router parameters. You can Telnet the Stack Master IP address:

1 Connect your station to the network.
2 Verify that you can communicate with the Avaya P330 using Ping to the IP of the Avaya P330. If there is no response using Ping, check the IP address and default gateway of both the Avaya P330 and the station.

Note: The Avaya P330 default IP address is 149.49.32.134 and the default subnet mask is 255.255.255.0.

3 From the Microsoft Windows® taskbar of your PC click **Start** and then **Run** (or from the DOS prompt of your PC), then start the Telnet session by typing:
telnet *<P330_IP_address>*
For example: **telnet 192.168.35.214.**
4 If the IP Address in Telnet command is the IP address of the stack, then connection is established with the Switch CLI entity of the Master module. When you see the "Welcome to P330" menu and are prompted for a Login Name, enter the default name **root**
5 When you are prompted for a password, enter the User Level password **root** in lower case letters (do NOT use uppercase letters). The User level prompt will appear when you have established communications with the Avaya P330.

Establishing a Modem (PPP) Connection with the P330

Overview

Point-to-Point Protocol (PPP) provides a Layer 2 method for transporting multi-protocol datagrams over modem links.

Connecting a Modem to the Console Port

A PPP connection with a modem can be established only after the Avaya P330 is configured with an IP address and net-mask, and the PPP parameters used in the Avaya P330 are compatible with the modem's PPP parameters.

1 Connect a terminal to the console port of the Avaya P330 switch as described in Connecting a Terminal to the Avaya P330 Serial port.
2 When you are prompted for a Login Name, enter the default name **`root`**.
3 When you are prompted for a password, enter the password **`root`**. You are now in Supervisor Level.
4 At the prompt, type:
`set interface ppp <``ip_addr><net-mask>`
with an IP address and netmask to be used by the Avaya P330 to connect via its PPP interface.

Note: The PPP interface configured with the `set interface ppp` command must be on a different subnet from the stack inband interface.

5 Set the baud rate, ppp authentication, and ppp time out required to match your modem. These commands are described in the "Command Line Interface" chapter.
6 At the prompt, type:
`set interface ppp enable`
The CLI responds with the following:
```
Entering the Modem mode within 60 seconds...
Please check that the proprietary modem cable is plugged
into the console port
```
7 Use the DB-25 to RJ-45 connector to plug the console cable to the modem's DB-25 connector. Plug the other end of the cable RJ-45 connector to the Avaya P330 console's RJ-45 port.
8 The Avaya P330 enters modem mode.
9 You can now dial into the switch from a remote station, and open a Telnet session to the PPP interface IP address.

Chapter 8

User Authentication

Introduction

A secure system provides safeguards to insure that only authorized personnel can perform configuration procedures. In Avaya P330, these safeguards form part of the CLI architecture and conventions.

CLI Architecture

The P330 stack supports both Layer 2 switching and Layer 3 switching.
The P333R-LB CLI includes two CLI entities to support this functionality.

- The Switch CLI entity is used to manage Layer 2 switching of the entire stack. CLI commands for managing Layer 2 switching are described in the *Avaya P330: Reference Guide*.
- The Router CLI entity is used to manage Layer 3 switching of a single module. CLI commands for managing Layer 3 switching are described in the *Avaya P330: Reference Guide*.

To switch between the entities, use the **`session`** command.

Security Levels

There are four security access levels – User, Privileged, Configure and Supervisor.

- The User level ('read-only') is a general access level used to show system parameter values.
- The Privileged level ('read-write') is used by site personnel to access stack configuration options.
- The Configure level is used by site personnel for Layer 3 configuration.
- The Supervisor level ('administrator') is used to define user names, passwords, and access levels of up to 10 local users. In Supervisor level you can also access RADIUS authentication configuration commands.

Note: If you wish to define more than ten users per switch, or accounts for a user on multiple switches, you should use RADIUS (Remote Authentication Dial-In User Service).

A login name and password are always required to access the CLI and the commands. The login name, password, and access-type (i.e., security level) for a

user account are established using the `username` command.

Switching between the entities, does not effect the security level since security levels are established specifically for each user. For example, if the operator with a privileged security level in the Switch entity switches to the Router entity the privileged security level is retained.

Note: If you wish to increase security, you can change the default user accounts and SNMP communities.

Note: The Web management passwords are the same as those of the CLI. If you change the passwords of the CLI then those passwords become active for Web management as well.

Entering the Supervisor Level

The Supervisor level is the level in which you first enter P330 CLI and establish user names for up to 10 local users. When you enter the Supervisor level, you are asked for a Login name. Type `root` as the Login name and the default password `root` (in lowercase letters):

```
Welcome to P330
Login: root
Password:****
Password accepted.
Cajun_P330-N(super)#
```

Defining new local users

Define new users and access levels using the following command in Supervisor Level.

In order to...	Use the following command...
Add a local user account and configure a user (name, password and access level)	username
To remove a local user account	no username

In order to...	Use the following command...
Display the username, password and access type for all users on the switch	show username

Exiting the Supervisor Level

To exit the Supervisor level, type the command `exit`.

Entering the CLI

To enter the CLI, enter your username and password. Your access level is indicated in the prompt as follows:

The User level prompt is shown below:

```
Cajun_P330-N>
```

The Privileged level prompt is shown below:

```
Cajun_P330-N#
```

The Configure level prompt for Layer 3 configuration is shown below:

```
P330-N(configure)#
```

The Supervisor level prompt is shown below:

```
Cajun_P330-N(super)#
```

RADIUS

Introduction to RADIUS

User accounts are typically maintained locally on the switch. Therefore, if a site contains multiple Avaya Switches, it is necessary to configure each switch with its own user accounts. Additionally, if for example a 'read-write' user has to be changed into a 'read-only' user, you must change all the 'read-write' passwords configured locally in every switch, in order to prevent him from accessing this level. This is obviously not effective management. A better solution is to have all of the user login information kept in a central location where all the switches can access it. P330 features such a solution: the Remote Authentication Dial-In User Service (RADIUS).

A RADIUS authentication server is installed on a central computer at the customer's site. On this server user authentication (account) information is configured that provides various degrees of access to the switch. The P330 will run as a RADIUS client. When a user attempts to log into the switch, if there is no local user account for the entered user name and password, then the switch will send an

Authentication Request to the RADIUS server in an attempt to authenticate the user remotely. If the user name and password are authenticated, then the RADIUS server responds to the switch with an Authentication Acknowledgement that includes information on the user's privileges ('administrator', 'read-write', or 'read-only'), and the user is allowed to gain access to the switch. If the user is not authenticated, then an Authentication Reject is sent to the switch and the user is not allowed access to the switch's embedded management.

The Remote Authentication Dial-In User Service (RADIUS) is an IETF standard (RFC 2138) client/server security protocol. Security and login information is stored in a central location known as the RADIUS server. RADIUS clients such as the P330, communicate with the RADIUS server to authenticate users.

All transactions between the RADIUS client and server are authenticated through the use of a “shared secret” which is not sent over the network. The shared secret is an authentication password configured on both the RADIUS client and its RADIUS servers. The shared secret is stored as clear text in the client’s file on the RADIUS server, and in the non-volatile memory of the P330. In addition, user passwords are sent between the client and server are encrypted for increased security.

Figure 8.1 illustrates the RADIUS authentication procedure:

Figure 8.1 RADIUS Authentication Procedure

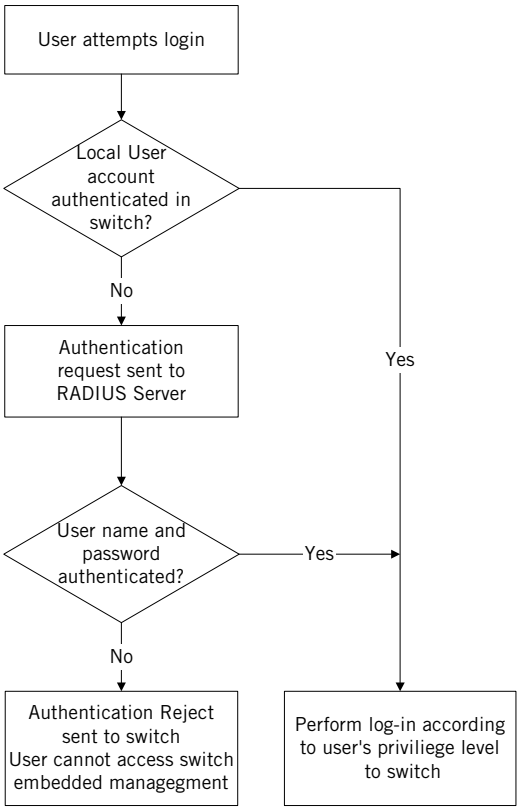
User attempts login
Local User account authenticated in switch?
No
Yes
Authentication request sent to RADIUS Server
User name and password authenticated?
Yes
No
Authentication Reject sent to switch
User cannot access switch embedded managegment
Perform log-in according to user's priviliege level to switch

Radius Commands

The following radius commands are accessible from Supervisor level.

In order to...	Use the following command...
Enable or disable authentication for the P330 switch. RADIUS authentication is disabled by default	set radius authentication
Set a primary or secondary RADIUS server IP address	set radius authentication server
Configure a character string to be used as a “shared secret” between the switch and the RADIUS server.	set radius authentication secret
Set the RFC 2138 approved UDP port number.	set radius authentication udp-port
Set the number of times an access request is sent when there is no response	set radius authentication retry-number
Set the time to wait before re-sending an access request.	set radius authentication retry-time
Remove a primary or secondary RADIUS authentication server	clear radius authentication server
Display all RADIUS authentication configurations. The shared secrets will not be displayed	show radius authentication

For a complete description of the RADIUS CLI commands, including syntax and output examples, refer to *Avaya P330: Reference Guide*.

Allowed Managers

With the Allowed Managers feature, the network manager can determine who may or may not gain management access to the switch. The feature can be enabled or disabled (default is disabled). When enabled, only those users that are configured in the Allowed Managers table are able to gain Telnet, HTTP, and SNMP management access to the switch.

You can configure up to 20 Allowed Mangers by adding or removing their IP address from the Allowed Managers List.

Note: The identification of an "Allowed Manager" is done by checking the Source IP address of the packets, thus if the Source IP address is modified on the way (NAT, Proxy, etc.), even an "Allowed Manager" will not be able to access the P330.

Allowed Manager CLI Commands

In order to...	Use the following command...
When set to enabled - only managers with ip address specified in the allowed table will be able to access the device	set allowed managers
Add/delete ip address of manager to/from the allowed table	set allowed managers ip
Show the IP addresses of the managers that are allowed to access the device	show allowed managers table
Show whether the status of allowed managers is enabled or disabled	show allowed managers status
Show the IP addresses of the managers that are currently connected	show secure current

AVAYA P333R-LB

SECTION 3: CONFIGURATION OF THE P330

Chapter 9

Basic Switch Configuration

Introduction

This chapter describes the parameters you can define for the chassis, such as its name and location, time parameters, and so on.

Use the CLI commands briefly described below for configuring the display on your terminal or workstation.

In order to...	Use the following command...
Open a CLI session to a P330 module in the stack, ATM or WAN expansion modules and Media Gateway Processor of G700.	session
Display or set the terminal width (in characters)	terminal width
Display or set the terminal length (in lines)	terminal length
Display or set the prompt	hostname
Return the prompt to its default value	no hostname
Clear the current terminal display	clear screen
Set the number of minutes before an inactive CLI session automatically logs out	set logout
Display the number of minutes before an inactive CLI session automatically times out	show logout
Access Layer 3 configuration if not logged in as supervisor (see “User Authentication“ chapter)	configure

System Parameter Configuration

Identifying the system

In order to make a P330 switch easier to identify, you can define a name for the switch, contact information for the switch technician and the location of the switch in the organization.

In order to...	Use the following command...
Configure the system name.	set system name
Configure the system contact person	set system contact
Configure the system location	set system location

Operating parameters

You can use the following commands to configure and display the mode of operation for the switch and display key parameters.

In order to...	Use the following command...
Configure the basic mode of operation of a module to either Layer 2 or Router	set device-mode
Display the mode of operation	show device-mode
Display system parameters	show system
Display module information for all modules within the stack	show module

Network Time Acquiring Protocols Parameter Configuration

The P330 can acquire the time form a Network Time Server. P330 supports the SNTP Protocol (RFC 958) over UDP port 123 or TIME protocol over UDP port 37. Use the CLI commands briefly described below for configuring and display time information and acquiring parameters.

In order to...	Use the following command...
Restore the time zone to its default, UTC.	clear timezone
Configure the time zone for the system	set timezone
Configure the time protocol for use in the system	set time protocol
Enable or disable the time client	set time client
Configure the network time server IP address	set time server
Display the current time	show time
Display the time status and parameters	show time parameters
Display the current time zone offset	show timezone
Get the time from the time server	get time

Default Settings of the P330

This section describes the procedures for the first-time configuration of the Avaya P330. The factory defaults are set out in detail in the tables included in this chapter.

Configuring the Switch

The Avaya P330 may be configured using the text-based Command Line Interface (CLI), the built-in Avaya P330 Device Manager (Embedded Web) or Avaya Multi-Service Network Manager™.

For instructions on the text-based CLI, see the *Avaya P330 Reference Guide*.

For instructions on installation of the graphical user interfaces, see Embedded Web Manager. For instructions on the use of the graphical user interfaces, refer to the Device Manager User's Guide on the Documentation and Utilities CD.

Avaya P330 Default Settings

The default settings for the Avaya P330 switch and its ports are determined by the Avaya P330 software. These default settings are subject to change in newer versions of the Avaya P330 software. See the Release Notes for the most up-to-date settings.

Table 10.1 Default Switch Settings

Function	Default Setting
IP address	149.49.32.134
Subnet Mask	255.255.255.0
Default gateway	0.0.0.0
Management VLAN ID	1
Spanning tree	Enabled
Bridge priority for Spanning Tree	32768
Keep alive frame transmission	Enabled
Network time acquisition	Enabled, Time protocol
Time server IP address	0.0.0.0

Table 10.1 Default Switch Settings

Function	Default Setting
Timezone offset	0 hours
SNMP communities: Read-only Read-write Trap SNMP	 Public Public Public
SNMP retries number	3
SNMP timeout	2000 Seconds
SNMP authentication trap	Disabled
CLI timeout	15 Minutes
User Name/Password	root/root

Table 10.2 Default Port Settings

Function	Default Setting		
	10/100Base-TX ports	**100Base-F ports**	**1000 Base-X ports**
Duplex mode	Full duplex	Full duplex	Full duplex only
Port Speed	100M	100M	1000M
Flow control	Off	Off	Off
Flow control advertisement	Off	N/A	Off (No pause)
Backpressure	On (only in Half duplex)	Not Applicable	Not Applicable
Autopartitioning	Disabled (only in Half duplex)	N/A	N/A
Auto-negotiation	Enable	Not Applicable	Enable[1]
Administration status	Enable	Enable	Enable
Port VLAN	1	1	1
Tagging mode	Clear	Clear	Clear

Table 10.2 Default Port Settings

Function	Default Setting		
Port priority	0	0	0
Spanning Tree cost	20	20	4
Spanning Tree port priority	128	128	128

1 Ensure that the other side is also set to Autonegotiation Enabled

Note: Functions operate in their default settings unless configured otherwise.

Chapter 11

Avaya P330 Layer 2 Features

This section describes the Avaya P330 Layer 2 features. It provides the basic procedures for configuring the P330 for Layer 2 operation.

Overview

The P330 family supports a range of Layer 2 features. Each feature has CLI commands associated with it. These commands are used to configure, operate, or monitor switch activity for each of the Layer 2 features.

This section of the *User's Guide* explains each of the features. Specifically, the topics discussed here include:

- Ethernet
- VLAN
- Port Based Network Access Control
- Spanning Tree Protocol
- Rapid Spanning Tree Protocol
- MAC Security
- Link Aggregation Group (LAG)
- Port Redundancy
- IP Multicast Filtering
- Stack Health
- Stack Redundancy
- Port Classification

Ethernet

Ethernet is one of the most widely implemented LAN standards. It uses the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access method to handle simultaneous demands. CSMA/CD is a multi-user network allocation procedure in which every station can receive the transmissions of every other station. Each station waits for the network to be idle before transmitting and each station can detect collisions by other stations.

The first version of Ethernet supported data transfer rates of 10 Mbps, and is therefore known as 10BASE-T.

Fast Ethernet

Fast Ethernet is a newer version of Ethernet, supporting data transfer rates of 100 Mbps. Fast Ethernet is sufficiently similar to Ethernet to support the use of most existing Ethernet applications and network management tools. Fast Ethernet is also known as 100BASE-T (over copper) or 100BASE-FX (over fiber).
Fast Ethernet is standardized as IEEE 802.3u.

Gigabit Ethernet

Gigabit Ethernet supports data rates of 1 Gbps. It is also known as 1000BASE-T (over copper) or 1000BASE-FX (over fiber).
Gigabit Ethernet is standardized as IEEE 802.3z.

Configuring Ethernet Parameters

Auto-negotiation

Auto-Negotiation is a protocol that runs between two stations, two switchs or a station and a switch. When enabled, Auto-Negotiation negotiates port speed and duplex mode by detecting the highest common denominator port connection for the endstations. For example, if one workstation supports both 10 Mbps and 100 Mbps speed ports, while the other workstation only supports 10 Mbps, then Auto-Negotiation sets the port speed to 10 Mbps.
For Gigabit ports, Auto-Negotiation determines the Flow Control configuration of the port.

Full-Duplex/Half-Duplex

Devices that support Full-Duplex can transmit and receive data simultaneously, as opposed to half-duplex transmission where each device can only communicate in turn.
Full-Duplex provides higher throughput than half-duplex.

Speed

The IEEE defines three standard speeds for Ethernet: 10, 100 and 1000 Mbps (also known as Ethernet, Fast Ethernet and Gigabit Ethernet respectively).

Flow Control

The process of adjusting the flow of data from one device to another to ensure that the receiving device can handle all of the incoming data. This is particularly important where the sending device is capable of sending data much faster than the receiving device can receive it.

There are many flow control mechanisms. One of the most common flow control protocols, used in Ethernet full-duplex, is called xon-xoff. In this case, the receiving device sends a an xoff message to the sending device when its buffer is full. The sending device then stops sending data. When the receiving device is ready to receive more data, it sends an *xon* signal.

Priority

By its nature, network traffic varies greatly over time, so short-term peak loads may exceed the switch capacity. When this occurs, the switch must buffer frames until there is enough capacity to forward them to the appropriate ports.

This, however, can interrupt time-sensitive traffic streams, such as Voice and other converged applications. These packets need to be forwarded with the minimum of delay or buffering. In other words, they need to be given high priority over other types of networkl traffic.

Priority determines in which order packets are sent on the network and is a key part of QoS (Quality of Service). The IEEE standard for priority on Ethernet networks is 802.1p.

Avaya P330 switches supports two internal priority queues – the High Priority queue and the Normal Priority queue.

- Packets tagged with priorities 4-7 are mapped to the High Priority queue; packets tagged with priorities 0-3 are mapped to the Normal Priority queue. This classification is based either on the packet's original priority tag, or, if the packet arrives at the port untagged, based on the priority configured for the ingress port (set using the `set port level` CLI command).

In cases where the packet was received tagged, this priority tag is retained when the packet is transmitted through a tagging port.

In cases where the priority is assigned based on the ingress priority of the port, then on an egress tagging port the packet will carry either priority 0 or priority 4, depending on the queue it was assigned to (High Priority=4, Normal Priority=0).

MAC Address

The MAC address is a unique 48-bit value associated with any network adapter. MAC addresses are also known as hardware addresses or physical addresses. They uniquely identify an adapter on a LAN.

MAC addresses are 12-digit hexadecimal numbers (48 bits in length). By convention, MAC addresses are usually written in one of the following two formats:

- MM:MM:MM:SS:SS:SS

- MM-MM-MM-SS-SS-SS

The first half of a MAC address contains the ID number of the device manufacturer. These IDs are regulated by an Internet standards body. The second half of a MAC address represents the serial number assigned to the device by the manufacturer.

CAM Table

The *CAM Table* contains a mapping of learned MAC addresses to port and VLANs. The switch checks forwarding requests against the addresses contained in the CAM Table:

- If the MAC address appears in the CAM Table, the packet is forwarded to the appropriate port.
- If the MAC address does not appear in the CAM Table, or the MAC Address mapping has changed, the frame is duplicated and copied to all the ports. Once a reply is received, the CAM table is updated with the new address/VLAN port mapping.

Ethernet Configuration CLI Commands

The following table contains a list of the configuration CLI commands for the Ethernet feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

Table 11.1 Configuration CLI Commands for Ethernet Feature

In order to...	Use the following command...
Set the auto negotiation mode of a port	set port negotiation
Administratively enable a port	set port enable
Administratively disable a port	set port disable
Set the speed for a 10/100 port	set port speed
Configure the duplex mode of a 10/100BASE-T port	set port duplex
Configure a name for a port	set port name
Set the send/receive mode for flow-control frames for a full duplex port	set port flowcontrol
Set the flow control advertisement for a Gigabit port when performing autonegotiation	set port auto-negotiation-flowcontrol-advertisement

In order to...	Use the following command...
Set the priority level of a port	set port level
Display settings and status for all ports	show port
Display per-port status information related to flow control	show port flowcontrol
Display the flow control advertisement for a Gigabit port used to perform auto-negotiation	show port auto-negotiation-flowcontrol-advertisement
Display the CAM table entries for a specific port	show cam
Clear all the CAM entries.	clear cam
Display the autopartition settings	show autopartition

Ethernet Implementation in the Avaya P333R-LB

This section describes the implementation of the Ethernet feature in the Avaya **P333R-LB**:

- Speed — 10/100 and 1G ports
- Priority queuing — 2 queues
- CAM size — 4K addresses

VLAN Configuration

VLAN Overview

A VLAN is made up of a group of devices on one or more LANs that are configured so that they operate as if they form an independent LAN, when in fact they may be located on a number of different LAN segments. VLANs can be used to group together departments and other logical groups, thereby reducing network traffic flow and increasing security within the VLAN.

The figure below illustrates how a simple VLAN can connect several endpoints in different locations and attached to different hubs. In this example, the Management VLAN consists of stations on numerous floors of the building and which are connected to both Device A and Device B.

Figure 11.1 VLAN Overview

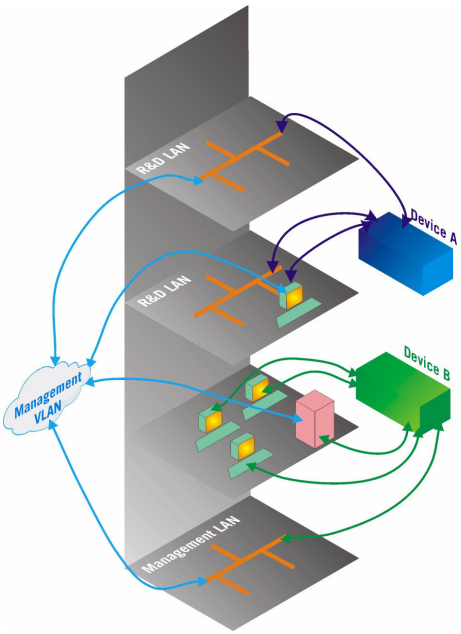

In virtual topological networks, the network devices may be located in diverse places around the LAN—such as in different departments, on different floors or in different buildings. Connections are made through software. Each network device is connected to a hub, and the network manager uses management software to assign each device to a virtual topological network. Elements can be combined into a VLAN even if they are connected to different devices.

VLANs should be used whenever there are one or more groups of network users that you want to separate from the rest of the network.

In Figure 11.2, the switch has three separate VLANs: Sales, Engineering, and

Marketing (Mktg). Each VLAN has several physical ports assigned to it with PC's connected to those ports. When traffic flows from a PC on the Sales VLAN for example, that traffic is *only* forwarded out the other ports assigned to that VLAN. Thus, the Engineering and Mktg VLANs are not burdened with processing that traffic.

Figure 11.2 VLAN Switching and Bridging

VLAN Tagging

VLAN Tagging is a method of controlling the distribution of information on the network. The ports on devices supporting VLAN Tagging are configured with the following parameters:

- Port VLAN ID
- Tagging Mode

The Port VLAN ID is the number of the VLAN to which the port is assigned. Untagged frames (and frames tagged with VLAN 0) entering the port are assigned the port's VLAN ID. Tagged frames are unaffected by the port's VLAN ID.

The Tagging Mode determines the behavior of the port that processes outgoing frames:

- If Tagging Mode is set to "Clear", the port transmits frames that belong to the port's VLAN table. These frames leave the device untagged.
- If Tagging Mode is set to "IEEE-802.1Q", all frames keep their tags when they leave the device. Frames that enter the switch without a VLAN tag will be tagged with the VLAN ID of the port they entered through.

Multi VLAN Binding

Multi VLAN binding (Multiple VLANs per port) allows access to shared resources by stations that belong to different VLANs through the same port. This is useful in applications such as multi-tenant networks, where each user has his a VLAN for privacy, but the whole building has a shared high-speed connection to the ISP. In order to accomplish this, P330 allows you to set multiple VLANs per port. The

three available Port Multi-VLAN binding modes are:

- **Bind to All** - the port is programmed to support the entire 3K VLANs range. Traffic from any VLAN is forwarded through a port defined as “Bind to All”. This is intended mainly for easy backbone link configuration
- **Bind to Configured** - the port supports all the VLANs configured in the switch/stack. These may be either Port VLAN IDs (PVID) or VLANs that were manually added to the switch.
- **Statically Bound** - the port supports VLANs manually configured on it.

Figure 11.3 illustrates these binding modes in P330.

Figure 11.3 Multiple VLAN Per-port Binding Modes

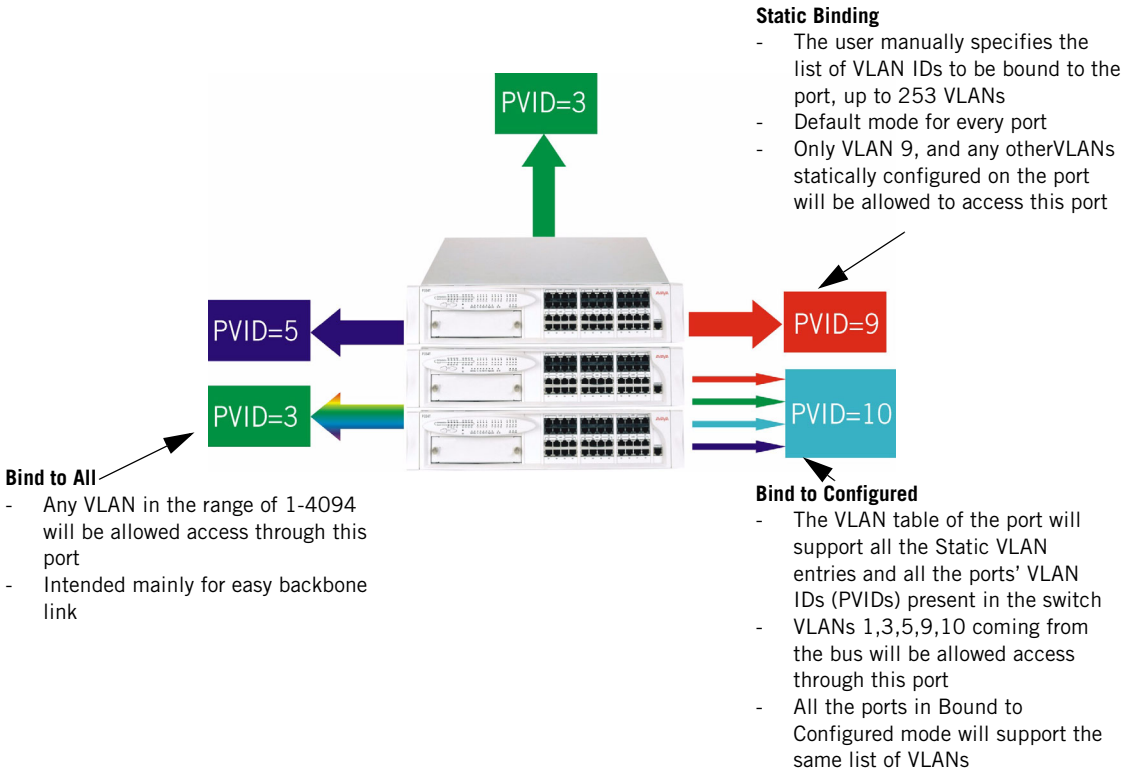
Static Binding
- The user manually specifies the list of VLAN IDs to be bound to the port, up to 253 VLANs
- Default mode for every port
- Only VLAN 9, and any otherVLANs statically configured on the port will be allowed to access this port
PVID=3
PVID=5
PVID=3
PVID=9
PVID=10
Bind to All
- Any VLAN in the range of 1-4094 will be allowed access through this port
- Intended mainly for easy backbone link
Bind to Configured
- The VLAN table of the port will support all the Static VLAN entries and all the ports' VLAN IDs (PVIDs) present in the switch
- VLANs 1,3,5,9,10 coming from the bus will be allowed access through this port
- All the ports in Bound to Configured mode will support the same list of VLANs

Ingress VLAN Security

When a VLAN-tagged packet arrives at a port, only the packets with the VLAN tag corresponding to the VLANs which are configured on the port will be accepted. Packets with other VLAN tags will be dropped.

VLAN CLI Commands

The following table contains a list of the CLI commands for the VLAN feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

Table 11.2 VLAN CLI Commands

In order to...	Use the following command...
Assign the Port VLAN ID (PVID)	set port vlan
Define the port binding method	set port vlan-binding-mode
Define a static VLAN for a port	set port static-vlan
Configure the tagging mode of a port	set trunk
Create VLANs	set vlan
Display the port VLAN binding mode settings	show port vlan-binding-mode
Display VLAN tagging information of the ports, port binding mode, port VLAN ID and the allowed VLANs on a port	show trunk
Display the VLANs configured in the switch.	show vlan
Clear VLAN entries	clear vlan
Clear a VLAN statically configured on a port	clear port static-vlan

VLAN Implementation in the Avaya P333R-LB

This section describes the implementation of the VLAN feature in the Avaya **P333R-LB**:

- No. of VLANs — 1024 tagged VLANs ranging from 1 to 3071

Port Based Network Access Control (PBNAC)

Port Based Network Access Control (IEEE 802.1X) is a method for performing authentication to obtain access to IEEE 802 LANs. The protocol defines an interaction between 3 entitites:

- Supplicant — an entity at one end of a point-to-point LAN segment that is being authenticated by an authenticator attached to the other end of that link.
- Authenticator — an entity at one end of a point-to-point LAN segment that facilitates authentication of the entity attached to the other end of that link; in this case, the P330.
- Authentication (RADIUS) Server — an entity that provides an authentication service to an authenticator. This service determines, from the credentials provided by the supplicant, whether the supplicant is authorized to access the services provided by the authenticator.

The process begins with the supplicant trying to access a certain restricted network resource, and upon successful authentication by the authentication server, the supplicant is granted access to the network resources.

How "Port Based" Authentication Works

802.1X provides a means of authenticating and authorizing users attached to a LAN port and of preventing access to that port in cases wher the authentication process fails. The authentication procedure is port based, which means:

- access control is achieved by enforcing authetication on connected ports
- if an end-point station that connects to a port is not authorized, the port state is set to "unauthorized" which closes the port to any traffic.
- As a result of an authentication attempt, the P330 port can be either in a "blocked" or a "forwarding" state.

802.1X interacts with existing standards to perform its authentication operation. Specifically, it makes use of Extensible Authentication Protocol (EAP) messages encapsulated within Ethernet frames (EAPOL), and EAP over RADIUS for the communication between the Authenticator and the Authentication Server.

PBNAC Implementation in the P330 Family

This section lists the conditions that govern the implementation of the 802.1X standard in the P330 line:

- You can configure PBNAC on the 10/100 Mbps Ethernet ports only.
- PBNAC can work only if a RADIUS server is configured on the P330 and the RADIUS server is carefully configured to support 802.1X.
- PBNAC and port/intermodule redundancy can co-exist on the same ports.
- PBNAC and LAGs can coexist on the same ports.
- PBNAC and Spanning Tree can be simultaneously active on a module.

Note: If either PBNAC or STP/RSTP are in a blocking state, the final state of the port will be blocked.

- When PBNAC is activated, the application immediately places all ports in a blocking state unless they were declared "Force Authenticate". They will be reverted to “Forwarding” state only when the port is authorized by the RADIUS server.

Note: The actual state of ports configured as "Force Authenticate" is determined by the STA.

Configuring the P330 for PBNAC

This section lists the basic tasks required to configure a P330 stack for PBNAC. To configure P330 for PBNAC, do the following:

- Configure a RADIUS server on a network reachable from the P330:
  - — Create user names and passwords for allowed users.
  - — Make sure the EAP option is enabled on this server.
- Configure the P330 for RADIUS:
  - — Configure RADIUS parameters.
  - — Enable the RADIUS feature.
  - — Configure the port used to access the RADIUS server as “force-authorized.”
- Connect the Supplicant—i.e., Windows XP clients—directly to the P330.
- Verify that the dot1x port-control is in auto mode.
- Set the dot1x system-auth-config to enable; the authentication process starts:
  - — The supplicant is asked to supply a user name and password.
  - — If authentication is enabled on the port, the Authenticator initiates authentication when the link is up.
  - — Authentification Succeeds: after the authentication process completes, the supplicant will receive a Permit/Deny notification.
  - — Authentication Fails: authentication will fail when the Supplicant fails to respond to requests from the Authenticator, when management controls prevent the port from being authorized, when the link is down, or when the user supplied incorrect logon information.

PBNAC CLI Commands

The following table contains a list of the CLI commands for the PBNAC feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

PBNAC CLI Commands

In order to...	Use the following command...
Configure dot1x on a system	set dot1x
Disable dot1x on all ports and return to default values	clear dot1x config
Display the system dot1x capabilities, protocol version, and timer values	show dot1x
Display all the configurable values associated with the authenticator port access entity (PAE) and backend authenticator	show port dot1x
Display all the port dot1x statistics	show port dot1x statistics
Set the minimal idle time between authentication attempts	set dot1x quiet-period
Set the time interval between attempts to access the Authenticated Station	set dot1x tx-period
Set the server retransmission timeout period for all ports	set dot1x server-timeout
Set the authentication period (an idle time between re-authentication attempts)	set dot1x re-authperiod
Set the authenticator-to-supplicant retransmission timeout period (the time for the switch to wait for a reply from the Authenticated Station)	set dot1x supp-timeout
Set the max-req for all ports (the maximal number of times the port tries to retransmit requests to the Authenticated Station before the session is terminated)	set dot1x max-req

In order to...	Use the following command...
Globally enable/disable 802.1x	set dot1x system-auth-control enable/disable
Set dot1x control parameter per port	set port dot1x port-control
Initialize port dot1x	set port dot1x initialize
Set the port to re-authenticate	set port dot1x re-authenticate
Set dot1x re-authentication mode per port	set port dot1x re-authentication
Set the 802.1x quiet period per port	set port dot1x quiet-period
Set the transmit period per port (a time interval between attempts to access the Authenticated Station)	set port dot1x tx-period
Set the supp-timeout per port (a time for the port to wait for a reply from the Authenticated Station)	set port dot1x supp-timeout
Set the server-timeout per port (a time to wait for a reply from the Authentication Server)	set port dot1x server-timeout
Set the re-authentication period per port (an idle time between re-authentication attempts)	set port dot1x re-authperiod
Set the max-req per port (the maximal number of times the port tries to retransmit requests to the Authenticated Station before the session is terminated)	set port dot1x max-req

Spanning Tree Protocol

Overview

Avaya P330 devices support both common Spanning Tree protocol (802.1d) and the enhanced Rapid Spanning Tree protocol (802.1w). The 802.1w is a faster and more sophisticated version of the 802.1d (STP) standard. Spanning Tree makes it possible to recover connectivity after an outage within a minute or so. RSTP, with its "rapid" algorithm, can restore connectivity to a network where a backbone link has failed in much less time.

In order to configure the switch to either common Spanning Tree or Rapid Spanning Tree protocol, use the `set spantree version` command.

Spanning Tree Protocol

The Spanning Tree Algorithm ensures the existence of a loop-free topology in networks that contain parallel bridges. A loop occurs when there are alternate routes between hosts. If there is a loop in an extended network, bridges may forward traffic indefinitely, which can result in increased traffic and degradation in network performance.

The Spanning Tree Algorithm:

- Produces a logical tree topology out of any arrangement of bridges. The result is a single path between any two end stations on an extended network.
- Provides a high degree of fault tolerance. It allows the network to automatically reconfigure the spanning tree topology if there is a bridge or data-path failure.

The Spanning Tree Algorithm requires five values to derive the spanning tree topology. These are:

1 A multicast address specifying all bridges on the extended network. This address is media-dependent and is automatically determined by the software.
2 A network-unique identifier for each bridge on the extended network.
3 A unique identifier for each bridge/LAN interface (a port).
4 The relative priority of each port.
5 The cost of each port.

After these values are assigned, bridges multicast and process the formatted frames (called Bridge Protocol Data Units, or BPDUs) to derive a single, loop-free topology throughout the extended network. The bridges exchange BPDU frames quickly, minimizing the time that service is unavailable between hosts.

Spanning Tree per Port

The Spanning Tree can take up to 30 seconds to open traffic on a port. This delay can cause problems on ports carrying time-sensitive traffic. You can therefore enable/disable Spanning Tree in P330 on a per-port basis to minimize this effect.

Rapid Spanning Tree Protocol (RSTP)

About the 802.1w Standard

The enhanced feature set of the 802.1w standard includes:

- Bridge Protocol Data Unit (BPDU) type 2
- New port roles: Alternate port, Backup port
- Direct handshaking between adjacent bridges regarding a desired topology change (TC). This eliminates the need to wait for the timer to expire.
- Improvement in the time it takes to propagate TC information. Specifically, TC information does not have to be propagated all the way back to the Root Bridge (and back) to be changed.
- Origination of BPDUs on a port-by-port basis.

Port Roles

At the center of RSTP—specifically as an improvement over STP (802.1d)—are the roles that are assigned to the ports. There are four port roles:

- Root port — port closest to the root bridge
- Designated port — corresponding port on the remote bridge of the local root port
- Alternate port — an alternate route to the root
- Backup port — an alternate route to the network segment

The RSTP algorithm makes it possible to change port roles rapidly through its fast topology change propagation mechanism. For example, a port in the “blocking” state can be assigned the role of “alternate port.” When the backbone of the network fails the port may be rapidly changed to forwarding.

Whereas the STA *passively* waited for the network to converge before turning a port into the forwarding state, RSTP *actively* confirms that a port can safely transition to forwarding without relying on any specific, programmed timer configuration.

RSTP provides a means of fast network convergence after a topology change. It does this by assigning different treatments to different port types. The port types and the treatment they receive follow:

- Edge ports — Setting a port to "edge-port" admin state indicates that this port is connected directly to end stations that cannot create bridging loops in the network. These ports transition quickly to forwarding state. However, if BPDUs are received on an Edge port, it’s operational state will be changed to "non-edge-port" and bridging loops will be avoided by the RSTP algorithm. The default admin state of all ports is "edge-port".

Note: You must manually configure uplink and backbone ports (including LAG logical ports) to be "non-edge" ports, using the CLI command `set port edge admin state`.

- Point-to-point Link ports — This port type applies only to ports interconnecting RSTP compliant switches and is used to define whether the devices are interconnected using shared Ethernet segment or pont-to-point Ethernet link. RSTP convergence is faster when switches are connected using point-to-point links. The default setting for all ports – automatic detection of point-to-point link – is sufficent for most networks.

Spanning Tree Implementation in the P330 Family

RSTP is implemented in P330 family of products so that it is interoperable with the existing implementation of STP. In order to configure the switch to either common Spanning Tree or Rapid Spanning Tree protocol, use the `set spantree version` command.

- After upgrading to software version 4.0, the default is spanning tree version STP. The default after NVRAM INIT remains STP.

The balance of this section lists the conditions and limitations that govern the implementation of Spanning Tree in the P330 line.

- RSTP's fast convergence benefits are lost when interacting with legacy (STP) bridges.
- When RSTP detects STP Bridge Protocol Data Units (BPDUs type 1) on a specific port, it will begin to "speak" 802.1d on this port only. Specifically, this means:
  - 802.1d bridges will ignore RSTP BPDUs and drop them.
  - 802.1d bridges will send 802.1d format BPDUs back to the switch.
  - The switch will change to 802.1d mode for that port only.

The P330 configured to RSTP is therefore able to simultaneously work with other switches implementing either RSTP or STP without specific user intervention.

- Spanning Tree configuration is performed on the stack level.
- If you do not upgrade all switches in the stack to firmware version 4.0, spanning tree will continue its normal operation. However, configuring Spanning Tree will not be possible until all switches are upgraded to version 4.0.
- RSTP is interoperable with P330 Port Redundancy and PBNAC applications. If either RSTP or PBNAC put the port in blocking, its final state will be "blocking".
- STP and Self Loop Discovery (SLD) are incompatible. However, If Spanning Tree is set to rapid-spanning-tree version, there is no need to use the Self-loop-discovery feature ; the RSTP algorithm avoids loops generated by the IBM token ring cabling.

- The 802.1w standard defines differently the default path cost for a port compared to STP (802.1d). In order to avoid network topology change when migrating to RSTP, the STP path cost is preserved when changing the spanning tree version to RSTP. You can use the default RSTP port cost by using the CLI command `set port spantree cost auto`.

Spanning Tree Protocol CLI Commands

The following table contains a list of CLI commands for the Spanning Tree feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

Table 11.3 Spanning Tree Protocol CLI Commandss

In order to...	Use the following command...
Enable/Disable the spanning tree application for the switch	set spantree
Set the bridge priority for spanning tree	set spantree priority
Set the RSTP bridge spanning tree max-age parameter	set spantree max-age
Set the RSTP bridge hello-time parameter	set spantree hello-time
Set the RSTP bridge forward-delay time prameter	set spantree forward-delay
Select between STP operation or RSTP switch operation	set spantree version
Display the bridge and per-port spanning tree information	show spantree
Set the TX hold count for the STA	set spantree priority
Add a port to the spanning tree application	set port spantree enable
Remove a port from the spanning tree application	set port spantree disable
Set the port spantree priority level	set port spantree priority
Set the cost of a port	set port spantree cost

Table 11.3 Spanning Tree Protocol CLI Commandss

In order to...	Use the following command...
Set the port as an RSTP port (and not as a common STA port)	set port spantree force-protocol-migration
Display a port's edge admin and operational RSTP state	show port edge state
Set the port as an RSTP edge port or non-edge port	set port edge admin state
Set the port point-to-point admin status	set port point-to-point admin status
Show the port's point-to-point admin and operational RSTP status	show port point-to-point status

MAC Security

The MAC security function is intended to filter incoming frames (from the line) with an unauthorized source MAC address (SA).

MAC Security Implementation in P330

When a frame is received on a secured port, its SA is checked against the MAC Address Table. If either the SA is not found there, or it is found but with a different port location, then the frame is rejected without being learned. A message is then sent to the CPU.

The Agent reports the attempted intrusion via an SNMP security violation trap containing the intruder's MAC address. To prevent the flooding of the Console's trap log / network, the Agent sends an intruder alert every 5 seconds for the first 3 times a specific intruder is detected on a port, and then every 15 minutes if the intrusion continues.

User should first enable the MAC security global mode (`set security mode`) and then configure the ports which should be secured (`set port security`). When setting a port to secured, the MAC addresses that a currently learnt on this port are preserved and considered as secure MAC, unless they are removed using `clear secure mac` command. Individual secure MACs can also be added.

Note: If the secure MAC editing command are to be implemented on a switch other than the stack master, a session should be opened to the relevant switch.

Note: Ports that are members of a port redundency scheme should not be also configured as secure ports.

MAC Security CLI Commands

The following table contains a list of the CLI commands for the MAC Security feature. The rules of syntax and output examples are all set out in detail in the *P330 Reference Guide*.

Table 11.4 MAC Security CLI Commands

In order to...	Use the following command...
Enable or disable the switch MAC security	set security mode

In order to...	Use the following command...
Enable or disable MAC security on a port	set port security
Add a unicast MAC address into the CAM table of a secured port (session command)	set secure mac
Remove a unicast MAC address from CAM table of a secured port (session command)	clear secure mac
Display the status of the MAC security feature (enabled/disabled)	show security mode
Display the secure MAC addresses of a port (session command)	show secure mac port
List the security mode of the ports of a switch	show port security

LAG

LAG Overview

A LAG uses multiple ports to create a high bandwidth connection with another device. For example: Assigning four 100BASE-T ports to a LAG on an Avaya P330 allows the switch to communicate at an effective rate of 400 Mbps with another switch.

LAGs provide a cost-effective method for creating a high bandwidth connection. LAGs also provide built-in redundancy for the ports that belong to a LAG. If a port in a LAG fails, its traffic is directed to another port within the LAG.

The behavior of the LAG is derived from the base port (the first port that becomes a LAG member). The attributes of the base port, such as port speed, VLAN number, etc., are applied to all the other member ports in the LAG.

When created, each LAG is automatically assigned a logical port number (usually designated 10x). This logical port number can then be used as any regular panel port for all configuration required for the LAG (Spanning Tree, Redundancy, etc.)

Note: In the P330-ML switches you need to erase **all** ports in t.he LAG in order to remove it.

LAG CLI Commands

The following table contains a list of the CLI commands for the LAG feature. The rules of syntax and output examples are all set out in detail in the *P330 Reference Guide*.

Table 11.5 LAG CLI Commands

In order to...	Use the following command...
Enable or disable a Link Aggregation Group (LAG) logical port on the switch	set port channel
Display Link Aggregation Group (LAG) information for a specific switch or port	show port channel

LAG Implementation in the Avaya P330 Family of Products

This section describes the implementation of the LAG feature in the P330 Family of products.

The P333R-LB supports up to 5 LAGs:

- Up to three LAGs from three groups of 8 10/100 Mbps ports:
  - Logical port 101 — ports1-4, 13-16
  - Logical port 102 — ports 5-8, 17-20
  - Logical port 103 — ports 9-12, 21-24
- Up to 2 LAGs (Logical ports 104-105) on the expansion module

Port Redundancy

Port redundancy involves the duplication of devices, services, or connections, so that, in the event of a failure, the redundant device, service, or connection can take over for the one that failed.

In addition to Link Aggregation Groups—which comprise the basic redundancy mechanism within the switch—the P330 offers an additional port redundancy scheme. To achieve port redundancy, you can define a redundancy relationship between any two ports in a stack. One port is defined as the primary port and the other as the secondary port. If the primary port fails, the secondary port takes over. You can configure up to 20 pairs of ports (or LAGs) per stack for port redundancy, and 1 pair per stack for intermodule redundancy. Each pair contains a primary and secondary port. You can configure any type of port to be redundant to any other.

Port Redundancy Operation

The Port Redundancy feature supports up to 20 pairs of ports per stack. The redundant or secondary port takes over when the primary port link is down. Port redundancy provides for the following in the P330:

- Switchback from the secondary to primary port is allowed
- Switching time intervals can be set by the user

Note: Port Redundancy interworks with the Spnning Tree Algorithm.

The Port Redundancy feature functions as follows:

- Port Redundancy enables the user to establish 20 pairs of ports. Within each pair, primary and secondary ports are defined. To prevent loops, only one port is enabled at a time.
- Following initialization, the primary port is enabled and the secondary port is disabled.
  - — If the active port link fails, the system enables the secondary port.
  - — If the secondary port is enabled and the primary port link becomes available again, the system will “switchback” to the primary port, unless configured otherwise by the user.
- Two timers are available:
  - — “min-time-between-switchovers” —minimum time (in seconds) between the failure of the primary port link and switchover to the secondary (backup) port.

Note: The first time the primary port fails, the switchover is immediate. This timer applies to subsequent failures.

— "switchback-interval" — the minimum time (in seconds) that the primary port link has to be up (following failure) before the system switches back to the primary port. The "none" parameter, if configured, prevents switching back to the primary.

Intermodule Port Redundancy

The intermodule port redundancy feature supports one pair of redundant ports per stack. The secondary port is activated:

- when the primary port link is down, or
- when the module in the stack holding the primary port has been powered down or removed.

Switching time for intermodule port redundancy is approximately 1 second.

Note: Defining intermodule port redundancy on ports with no link causes both ports to be disabled. You should connect the link prior to attempting to define intermodule port redundancy.

Note: Once a port has been designated in a redundancy scheme, either as a primary or a secondary port, it can not be designated in any other redundancy scheme.

Note: Intermodule Port Redundancy does not interworks with the Spnning Tree Algorithm.

Port Redundancy CLI Commands

The following table contains a list of the CLI commands for the Redundancy feature. The rules of syntax and output examples are all set out in detail in the *P330 Reference Guide*.

Table 11.6 Redundancy CLI Commands (check spec)

In order to...	Use the following command...
Define or remove port redundancy schemes	set port redundancy
Enable the defined port redundancy schemes	set port redundancy enable

In order to...	Use the following command...
Disable the defined port redundancy schemes	set port redundancy disable
Define the timers that control the port redundancy operation	set port redundancy-interval
Display information on port redundancy schemes.	show port redundancy
Define the switch's unique intermodule redundancy scheme	set intermodule port redundancy
Clear the intermodule redundancy	set intermodule port redundancy off
display the intermodule redundancy entry defined for the switch	show intermodule port redundancy

IP Multicast Filtering

Overview

IP Multicast is a method of sending a single copy of an IP packet to multiple destinations. It can be used by different applications including video streaming and video conferencing.

The Multicast packet is forwarded from the sender to the recipients, duplicated only when needed by routers along the way and sent in multiple directions such that it reaches all the members of the Multicast group. Multicast addresses are a special kind of IP addresses (class D), each identifying a multicast group. Stations join and leave multicast groups using IGMP. This is a control-plane protocol through which IP hosts register with their router to receive packets for certain multicast addresses.

IP multicast packets are transmitted on LANs in MAC multicast frames. Traditional LAN switches flood these multicast packets like broadcast packets to all stations in the VLAN. In order to avoid sending multicast packets where they are not required, multicast filtering functions may be added to the layer 2 switches, as described in IEEE standard 802.1D. Layer 2 switches capable of multicast filtering send the multicast packets only to ports connecting members of that multicast group. This is typically based on IGMP snooping.

The Avaya P330 supports multicast filtering. The P330 learns which switch ports need to receive which multicast packets and configures the necessary information into the switch's hardware tables. This learning is based on IGMP (version 1 or 2) snooping.

The multicast filtering function in the P330 is transparent to the IP hosts and routers. It does not affect the forwarding behavior apart from filtering multicast packets from certain ports where they are not needed. To the ports that do get the multicast, forwarding is performed in the same way as if there was no filtering, and the multicast packet will not be sent to any ports that would not receive it if there was no filtering.

The multicast filtering function operates per VLAN. A multicast packet arriving at the device on a certain VLAN will be forwarded only to a subset of the ports of that VLAN. If VLAN tagging mode is used on the output port, then the multicast packet will be tagged with the same VLAN number with which it arrived. This is interoperable with multicast routers that expect Layer 2 switching to be done independently for each VLAN.

IP Multicast Filtering configuration is associated with the setting up of three timers:

- The **Router Port Pruning** timer ages out Router port information if IGMP queries are not received within the configured time.
- The **Client Port Pruning** time is the time after the P330 switch reset that the filtering information is learned by the switch but not configured on the ports.
- The **Group Filtering Delay** time is the time that the switch waits between becoming aware of a Multicast group on a certain VLAN and starting to filter traffic for this group.

IP Multicast CLI Commands

The following table contains a list of the CLI commands for the IP Multicast feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

Table 11.7 IP Multicast CLI Commands

In order to...	Use the following command...
Enable or disable the IP multicast filtering application	set intelligent-multicast
Define aging time for client ports	set intelligent-multicast client port pruning time
Define aging time for router ports	set intelligent-multicast router port pruning time
Define group filtering time delays	set intelligent-multicast group-filtering delay time
Display the status IP multicast filtering application	show intelligent-multicast
Shows whether the connected unit's hardware supports IP multicast filtering	show intelligent-multicast hardware-support

IP Multicast Implementation in the Avaya P333R-LB

This section describes the implementation of the IP multicast feature in the Avaya **P333R-LB**:

- No. of multicast groups — 1000

Stack Health

The P330 software provides a Stack Helath feature for verifying the integrity of the P330 stack cascading module and cables.

Overview

The Stack Health feature will identify defective modules and cables that may be installed in the P330 stack. The Stack Health algorithm separately checks all stacking modules and the Octaplane connections (including Redundant cable).

Implementation of Stack Health in the P330 Family

When activating the Stack Health feature, the agents in all modules start sending special packets of various length via all stacking cables to one another. The Master module synchronizes this process and collects the results.

- When the Redundant Cable is present, the user is prompted to disconnect one of the short Octaplane cables and the redundant connection will be checked. Then, when prompted, the cable should be reconnected and the test will run a second time to check the regular Octaplane connections.
- The stack is reset after the Stack Health process completes.

Note: You should not load the stack with traffic during this test.

Note: If the stack health process fails, try to fasten or replace the stack cable between the modules where the failure has occurred. If the problem persists, try to fasten or replace either or both of the stacking modules.

Stack Health CLI Commands

The following table contains a list of the CLI commands for the Stack Health feature. The rules of syntax and output examples are all set out in detail in the *Reference Guide*.

Table 11.8 Stack Health CLI Command

In order to...	Use the following command...
Initiate the stack health testing procedure	set stack health

Port Classification

Overview

With the P330, you can classify any port as regular or valuable. Setting a port to valuable means that, in case of Ethernet link failure of that port, a link fault trap can be sent even when the port is disabled and a fast aging operation on the CAM table will be performed. This feature is particularly useful for the link/intermodule redundancy application, where you need to be informed about a link failure on the dormant port and resume traffic quickly.

Port Classification CLI Commands

In order to...	Use the following command...
Set the port classification to either regular or valuable	set port classification
Display a port's classification	show port classification

Stack Redundancy

In the unlikely event that a P330 switch or Octaplane link should fail, stack integrity is maintained if the redundant cable is connected to the stack. The broken link is bypassed and data transmission continues uninterrupted. The single management IP address for the stack is also preserved for uninterrupted management and monitoring. You can remove or replace any unit within the stack without disrupting operation or performing stack-level reconfiguration.

Since each P330 module has an integral SNMP agent, any module in a stack can serve as the stack Network Managment Agent (NMA) while other NMAs act as redundant agents in “hot” standby. If the “live” NMA fails then a backup is activated instantaneously.

Chapter 12

Avaya P330 Layer 3 Features

Introduction

This section describes the Avaya P330 Layer 3 features. It provides the basic procedures for configuring the P330 for Layer 3 operation.

What is Routing?

Routing enables transfer of a data packet from source to destination using a device called a router. Routing involves two basic activities:

- determining optimal routing paths
- transmitting information packets through an internetwork

Routers use routing tables to determine the routes to particular network destinations and, in some cases, metrics associated with those routes. Routers communicate with one another and maintain their routing tables through the transmission of a variety of messages.

The Routing Update Message is one such message. Routing Updates generally consist of all or a portion of a routing table. By analyzing Routing Updates from all routers, a router can build a detailed picture of network topology.

A Link-State Advertisement is another example of a message sent between routers. Link-State Advertisements inform other routers of the state of the sender's links. Link information can also be used to build a complete picture of the network's topology. Once the network topology is understood, routers can determine optimal routes to network destinations.

Routers can route only those messages that are transmitted in a routable protocol, such as IP or IPX. Messages in non-routable protocols, such as NetBIOS and LAT, cannot be routed, but they can be transferred from LAN to LAN via a bridge.

When a router receives a packet, it examines the packet's destination protocol address. The router then determines whether it knows how to forward the packet to the next hop. If the router does not know how to forward the packet, it typically drops the packet unless a default gateway is defined. If the router knows how to forward the packet, it changes the packet destination’s physical address to that of the next hop and transmits the packet.

The next hop may or may not be the ultimate destination host. If not, the next hop is usually another router, which executes the same switching decision process. As the packet moves through the internetwork, its physical address changes but its protocol address remains constant. This process is illustrated in the figure below.

Figure 12.1 Routing

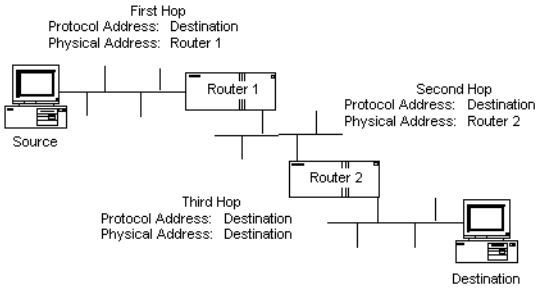

The relation between the destination host's protocol address and its physical address is obtained by the routers using the ARP request/reply mechanism, and the information is stored within the ARP table in the router (see "The ARP Table" on page 111).

Within an enterprise, routers serve as an intranet backbone that interconnects all networks. This architecture strings several routers together via a high-speed LAN topology such as Fast Ethernet or Gigabit Ethernet. Within the global Internet, routers do all the packet switching in the backbones.

Another approach within an enterprise is the collapsed backbone, which uses a single router with a high-speed backplane to connect the subnets, making network management simpler and improving performance.

Routing Configuration

Forwarding

The P330 forwards IP packets between IP networks. When it receives an IP packet through one of its interfaces, it forwards the packet through one of its interfaces. The P330 supports multinetting, enabling it to forward packets between IP subnets on the same VLAN as well as between different VLANs. Forwarding is performed through standard means in Router mode.

Multinetting (Multiple Subnets per VLAN)

In Router Mode, most applications such as RIP and OSPF, operate per IP interface. Other applications such as VRRP and DHCP/BOOTP Relay operate per VLAN. Configuration of these applications is done in the Interface mode. When there is only a single interface (subnet) per VLAN then system behavior is intuitive since a subnet and a VLAN are the same.

Multiple interfaces (subnets) per VLAN are more complicated. For example, if there are two interfaces over the same VLAN and you configure DHCP server on one interface, it will be used also for the second interface over the same VLAN. This behavior might be less expected and in some cases wrong.

In order to prevent misconfiguration and unexpected results, the P330 prevents configuration of VLAN-oriented commands on an interface unless explicitly requested by the user via the “enable vlan commands” CLI command.

Configuring the “enable vlan commands” on one interface defeats this option on other interfaces that belong to the same VLAN. This ensures that VLAN-oriented commands can be configured from one interface only.

In case there is only one interface over a VLAN, then VLAN oriented commands for this VLAN can be configured through the single interface without using the “enable vlan commands” command.

Note:
1. VLAN-oriented commands that were configured affect the VLAN of the interface that was used at the time the command was issued.
2. If the interface is moved to another VLAN (using the “ip vlan command”) VLAN oriented configuration still relates to the original VLAN.

IP Configuration

IP Configuration CLI Commands

In order to...	Use the following command...
Enable IP routing	ip routing
Set ICMP error messages	ip icmp-errors
Specify the format of netmasks in the show command output	ip netmask-format
Create and/or enter the Interface Configuration Mode	interface
Assign an IP address and mask to an interface	ip address
Set the administrative state of an IP interface	ip admin-state
Update the interface broadcast address	ip broadcast-address
Define a default gateway (router)	ip default-gateway
Define the interface RIP route metric value	default-metric
Enable net-directed broadcast forwarding	ip directed-broadcast
Set the IP routing mode of the interface	ip routing-mode
Enable or disable the sending of redirect messages on the interface	ip redirect
Check host reachability and network connectivity	ping
Trace route utility	traceroute
Create a router Layer 2 interface	set vlan (Layer 3)
Specify the VLAN on which an IP interface resides	ip vlan/ip vlan name

In order to...	Use the following command...
Use this command before configuring VLAN-oriented parameters, when there is more than one interface on the same VLAN	enable vlan commands
Display information about the IP unicast routing table	show ip route (Layer 3)
Display information for an IP interface	show ip interface
Display the status of ICMP error messages	show ip icmp
Delete a Layer 2 Router interface	clear vlan

Assigning Initial Router Parameters

To configure the initial router parameters perform the following via the CLI:

1 Enter **`set device-mode router`** and press Enter.
You will be prompted to reset the module.

2 Type **`y`**.
Wait for the module to restart and for the CLI prompt to reappear.

3 Type **`show device-mode`** and press Enter to ensure that the module is in router mode.

Note: Assign the stack IP address as described in Assigning P330's IP Stack Address before you assign the Initial Router IP address.

4 To access Router commands from the Master module, type the command **`session`** `<module number>` **`router`** where `<module number>` is the location of the router module in the stack, and press Enter.
The command prompt changes from `P330-N>` to `Router-N#>` where N is the number of the router in the stack (see P330 Sessions).

5 Type **`configure`** and press Enter. The prompt `Router-N(configure)#` appears.

Note: If the IP interface is on VLAN #1, continue with step 7.

6 Create the management/routing VLAN. Use the command **set vlan** <Vlan-id> **name** <Vlan-name> replacing <Vlan-id> by the VLAN number, and <Vlan-name> by the VLAN name. Press Enter.

7 Create an IP interface name. Type:
Router(configure)# interface <interface-name>
Press Enter.
The **Router(configure-if:<interface-name>)#** prompt appears.

8 Assign the IP address and network mask of the IP interface you have created. Use the command**:**
Router(configure-if:<interface-name>)# ip address <ip-address> <netmask>
Press Enter

9 Assign a vlan to the IP interface you have created. Type: Assign a vlan to the IP interface you have created. Type:
Router(configure)# interface <interface-name># ip vlan <vland-id>
Press Enter.

10 Type **exit** and press Enter. This returns you to the prompt:
Router(configure)#

11 If the management station is not on the same subnet as the switch, configure a default gateway (static route). Use the command:
ip default-gateway <ip-address> and press Enter, replacing <ip-address> with the IP address of the default gateway.

12 Save the configuration changes by typing **copy running-config startup-config** and press Enter.

RIP (Routing Interchange Protocol) Configuration

RIP Overview

RIP is a "distance vector protocol"— that is, the router decides which path to use on distance (the number of intermediate hops). In order for this protocol to work correctly, all of the routers (and possibly the nodes) need to gather information on how to reach each destination in the Internet. The very simplicity of RIP has a disadvantage, however: it does not take into account the network bandwidth, physical cost, data priority, and so on.

The P330 supports the widely used RIP routing protocol (both RIPv1 and RIPv2). The RIPv1 protocol imposes some limitations on the network design with regard to subnetting. When operating RIPv1, you must not configure variable length subnet masks (VLMS). Each IP network must have a single mask, implying that all subnets in a given IP network are of the same size. Also, when operating RIPv1, you must not configure supernets, which are networks with a mask smaller than the natural net mask of the address class, such as 192.1.0.0 with mask 255.255.0.0 (smaller than the natural class C mask which is 255.255.255.0). For detailed descriptions of RIP refer to the standards and published literature.

RIPv2 is a new version of the RIP routing protocol but with some advantages over RIPv1. RIPv2 solves some of the problems associated with RIPv1. The most important change in RIPv2 is the addition of a subnet mask field which allows RIPv2 to support variable length subnets. RIPv2 also includes an authentication mechanism similar to the one used in OSPF.

The RIP version, 1 or 2, is configured per IP interface. Configuration must be homogenous on all routers on each subnet—there can not be both RIPv1 and RIPv2 routers configured on the same subnet. However, different IP interfaces of the P330 can be configured with different RIP versions (as long as all routers on the subnet are configured to the same version).

RIPv2 and RIPv1 are considered the same protocol with regard to redistribution to/from OSPF and static route preferences.

RIP2

RIP2 overcomes some of the shortcomings of RIP. The table below summarizes the differences between RIP and RIP2.

Table 12.1 DIfferences Between RIP and RIP2

RIP2	RIP
Multicast addressing	Broadcast Addressing
Event-driven	Timer-based (update every 30 seconds)
VLSM support (subnet information transmitted)	Fixed subnet masks

RIP CLI Commands

In order to...	Use the following command...
Configure the Routing Information Protocol (RIP)	router rip
Specify a list of networks on which the RIP is running	network
Redistribute routing information from other protocols into RIP	redistribute
Specify the RIP version running on the interface basis	ip rip rip-version
Set the interface RIP route metric value	default-metric
Set the RIP Send and Receive mode on an interface	ip rip send-receive
Enable learning of the default route received by the RIP	ip rip default-route-mode
Enable split-horizon with poison-reverse on an interface	ip rip poison-reverse
Enable split-horizon mechanism	ip rip split-horizon

In order to...	Use the following command...
Specify the type of authentication used in RIP Version 2 packets	ip rip authentication mode
Set the authentication string used on the interface	ip rip authentication key
Specify the RIP timers values	timers basic

OSPF (Open Shortest Path First) Configuration

OSPF Overview

OSPF is a routing protocol developed for IP networks based on the shortest path first or link-state algorithm. It was introduced to overcome the limitations of RIP in increasingly complex network designs.

OSPF is based on the cost of a particular path. In contrast, RIP uses hops as a path criterion. Also, updates are sent on a "need to know" basis rather than every 30 seconds as with RIP.

The advantage of shortest path first algorithms is that they result in smaller more frequent updates everywhere. They converge quickly, thus preventing such problems as routing loops and Count-to-Infinity (when routers continuously increment the hop count to a particular network). This stabilizes the network.

The disadvantage of shortest path first algorithms is that they require a lot of CPU power and memory. .

Routers use link-state algorithms to send routing information to all nodes in an internetwork by calculating the shortest path to each node. This calculation is based on a topography of the Internet constructed by each node. Each router sends that portion of the routing table (keeps track of routes to particular network destinations) that describes the state of its own links, and it also sends the complete routing structure (topography).

The P330 supports the OSPF routing protocol. The P330 can serve as an OSPF Autonomous System Boundary Router (ASBR) by configuration of route redistribution. The P330 can be installed in the OSPF backbone area (area 0.0.0.0) or in any OSPF area that is part of a multiple areas network. However, the P330 cannot be configured to be an OSPF area border router itself.

The P330 supports the equal-cost multipath (ECMP) feature which allows load balancing by splitting traffic between several equivalent paths.

While OSPF can be activated with default values for each interface using a single command, many of the OSPF parameters are configurable.

For a detailed description of OSPF, refer to the OSPF standards and published literature.

OSPF CLI Commands

In order to...	Use the following command...
Enable OSPF protocol	router ospf
Configure the area ID of the router	area
Configure router identity	ip ospf router-id
Configure a passive ospf interface	passive-interface
Redistribute routing information from other protocols into OSPF	redistribute
Configure the delay between runs of OSPF's SPF calculation	timers spf
Configure interface metric	ip ospf cost
Specify the time interval between hellos the router sends	ip ospf hello-interval
Configure the interval before declaring the neighbor as dead.	ip ospf dead-interval
Configure interface priority used in DR election	ip ospf priority
Configure the interface authentication password	ip ospf authentication-key
Display general information about OSPF routing	show ip ospf
Display the OSPF-related interface information	show ip ospf interface
Display OSPF neighbor information on a per-interface basis	show ip ospf neighbor
Display lists of information related to the OSPF database for a specific router	show ip ospf database
Configure an interface as passive	passive-interface

Static Routing Configuration

Static Routing Overview

When dynamic routing protocols (RIP or OSPF) are not appropriate, you can manually configure *static routes* to indicate the next hop on the path to the final packet destination.

A static route becomes inactive if the interface over which it is defined is disabled. When the interface is enabled, the static route becomes active again. They are never timed-out, or lost over reboot, and can only be removed by manual configuration. Deletion (by configuration) of the IP interface deletes the static routes using this interface as well.

Static routes can only be configured for remote destinations, i.e. destinations that are reachable via another router as a next hop. The next hop router must belong to one of the directly attached networks for which the P330 has an IP interface. "Local" static routes, such as those that have no next hop, are not allowed.

Two kinds of static routes can be configured:

- High Preference static routes which are preferred to routes learned from any routing protocol
- Low Preference static routes which are used temporarily until the route is learned from a routing protocol. By default, a static route has Low Preference.

Static routes can be advertised by routing protocols (i.e., RIP, OSPF) as described under Route redistribution.

Static routes also support load-balancing similar to OSPF. A static route can be configured with multiple next hops so that traffic is split between these next hops.

This can be used for example to load-balance traffic between several firewalls which serve as the default gateway.

Static Routing Configuration CLI Commands

In order to...	Use the following command...
Establish a static route	ip route
Remove a static route	no ip route
Set the maximum number of route entries in the routing table	ip max-route-entries
Set the maximum number of route entries in the routing table to the default value	no ip max-route-entries

In order to...	Use the following command...
Define a default gateway (router)	ip default-gateway
Remove the default gateway (router)	no ip default-gateway
Delete all the dynamic routing entries from the Routing Table	clear ip route
Display information about the IP unicast routing table	show ip route
Display a routing table for a destination address	show ip route best-match
Display the static routes	show ip route static
Display the number of routes known to the switch	show ip route summary

Route Preferences

The routing table may contain routes from different sources. Routes to a certain destination may be learned independently from RIP and from OSPF, and at the same time, a static route can also be configured to the same destination. While metrics are used to choose between routes of the same protocol, protocol preferences are used to choose between routes of different protocols.

The preferences only apply to routes for the same destination IP address and mask. They do not override the longest-match choice. For example, a high-preference static default route will not be preferred over a RIP route to the subnet of the destination.

P330 protocol preferences are listed below from the most to the least preferred:

1 Local (directly attached net)
2 High-preference static (manually configured routes)
3 OSPF internal routes
4 RIP
5 OSPF external routes
6 Low-preference static (manually configured routes).

Route Redistribution

Route redistribution is the interaction of multiple routing protocols. OSPF and RIP can be operated concurrently in the P330. In this case, the P330 can be configured to redistribute routes learned from one protocol into the domain of the other routing protocol. Similarly, static routes may be redistributed to RIP and to OSPF. Route redistribution involves metric changes and sometimes causes routing loops in the presence of other routes with incompatible schemes for route redistribution and route preferences. Be careful!

The the P330 scheme for metric translation in route redistribution is as follows:

- Static to RIP metric configurable (default 1)
- OSPF internal metric N to RIP metric 1
- OSPF external type 1 metric N to RIP metric 1
- OSPF external type 2 metric N to RIP metric N+1
- Static to OSPF external type 2, metric configurable (default 1)
- RIP metric N to OSPF external type 2, metric N
- Direct to OSPF external type 2, metric 1.

By default, the P330 does not redistribute routes between OSPF and RIP. Redistribution from one protocol to the other can be configured. By default, static routes are not redistributed to RIP and OSPF. The P330 allows the user to globally enable redistribution of static routes to RIP, and separately, to globally enable redistribution of static routes to OSPF. In addition, the P330 lets the user configure, on a per static route basis, whether the route is to be redistributed to RIP and OSPF, and what metric (in the range of 1-15). The default state is to enable the route to be redistributed at metric 1. When static routes are redistributed to OSPF, they are always redistributed as external type 2.

Route Redistribution Commands

In order to...	Use the following command...
Redistribute routing information from other protocols	redistribute

ARP (Address Resolution Protocol) Table Configuration

ARP Overview

IP logical network addresses are independent of physical addresses. Since the physical address must be used to convey data in the form of a frame from one device to another, a mechanism is required to acquire a destination device hardware address from its IP address. This mechanism/ability is called ARP (Address Resolution Protocol).

The following mechanism describes how a station builds an ARP Table:

Figure 12.2 Building an ARP Table

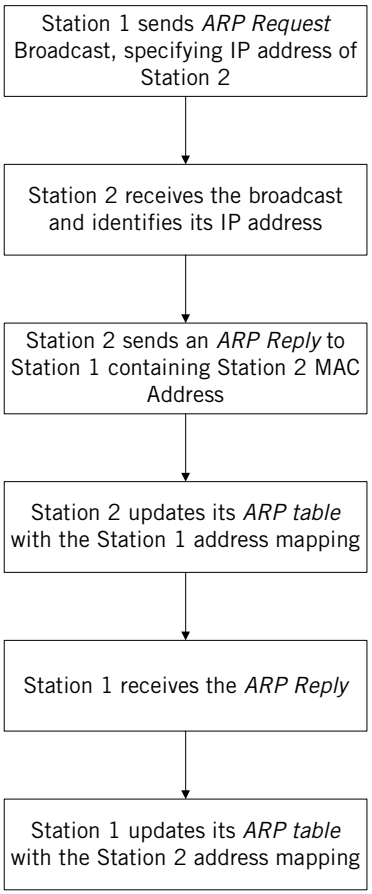
Station 1 sends *ARP Request* Broadcast, specifying IP address of Station 2
Station 2 receives the broadcast and identifies its IP address
Station 2 sends an *ARP Reply* to Station 1 containing Station 2 MAC Address
Station 2 updates its *ARP table* with the Station 1 address mapping
Station 1 receives the *ARP Reply*
Station 1 updates its *ARP table* with the Station 2 address mapping

The ARP Table

The ARP table is used to store recently used pairs of IP/MAC addresses. This saves time and communication costs, since the host looks in the ARP cache first when transmitting a packet. If the information is not there, then the host sends an ARP Request (see Figure 12.2).

ARP CLI Commands

In order to...	Use the following command...
Add a permanent entry to the Address Resolution Protocol (ARP) cache	arp
Configure the amount of time that an entry remains in the ARP cache	arp timeout
Set the amount of time that an entry remains in the ARP cache back to default	no arp timeout
Set the maximum number of ARP cache entries allowed in the ARP cache	ip max-arp-entries
Set the maximum number of ARP cache back to default	no ip max-arp-entries
Enable proxy ARP on an interface	ip proxy-arp
Disable proxy ARP on an interface	no ip proxy-arp
Delete all dynamic entries from the ARP cache and the IP route cache	clear arp-cache
Display the Address Resolution Protocol (ARP) cache	show ip arp
Display the IP address of a host, based on a known MAC address	show ip reverse-arp

BOOTP/DHCP (Dynamic Host Configuration Protocol) Relay Configuration

BOOTP/DHCP Overview

BOOTP

Short for Bootstrap Protocol, BootP is an Internet protocol that enables a diskless workstation to discover its own IP address, the IP address of a BOOTP server on the network, and a file to be loaded into memory to boot the machine. This enables the workstation to boot without requiring a hard or floppy disk drive. It is used when the user/station location changes frequently.

The protocol is defined by RFC 951.

DHCP

Short for Dynamic Host Configuration Protocol, DHCP assigns dynamic IP addresses to devices on a network. With dynamic addressing, a device can have a different IP address every time it connects to the network. In some systems, the device's IP address can even change while it is still connected. DHCP also supports a mix of static and dynamic IP addresses.

Dynamic addressing simplifies network administration because the software keeps track of IP addresses rather than requiring an administrator to manage the task. This means that a new computer can be added to a network without the hassle of manually assigning it a unique IP address. Many ISPs use dynamic IP addressing for dial-up users. However, dynamic addressing may not be desirable for a network server.

DHCP/BOOTP Relay

The P330 supports the DHCP/BOOTP Relay Agent function. This is an application that accepts DHCP/BOOTP requests that are broadcast on one VLAN and sends them to a DHCP/BOOTP server that connects to another VLAN or a server that may be located across one or more routers that would otherwise not get the broadcast request. The relay agent handles the DHCP/BOOTP replies as well, transmitting them to the client directly or as broadcast, according to a flag in the reply message. Note that the same DHCP/BOOTP relay agent serves both the BOOTP and DHCP protocols.

When there is more than one IP interface on a VLAN, the P330 automatically chooses one of the IP interface's to determine the relay network. Alternatively, you can configure the relay networks that the P330 will use. If you have defined more than one network, the P330 selects the network to be relayed on a Round Robin basis.

The DHCP/BOOTP server uses the relayed network information to decide from which subnet the address should be allocated. Therefore, the DHCP/BOOTP server must be configured to allocate addresses from the relayed networks configured on the P330.

DHCP/BOOTP Relay in P330 is configurable per VLAN and allows you to specify two DHCP/BOOTP servers. In this case, it duplicates each request, and sends it to both servers. This provides redundancy and prevents the failure of a single server from blocking hosts from loading.

You can enable or disable or DHCP/BOOTP Relay in P330.

BOOTP/DHCP CLI Commands

In order to...	Use the following command...
Enable or disable relaying of bootp and dhcp requests to the BOOTP/ DHCP server	[no] ip bootp-dhcp relay (no - restores default)
Add or remove a BOOTP/DHCP server to handle BOOTP/DHCP requests received by this interface	[no] ip bootp-dhcp server (no - restores default)
Select the networks from which the bootp/dhcp server shall allocate an address	[no] ip bootp-dhcp network (no - restores default)

NetBIOS Re-broadcast Configuration

NetBIOS Overview

Short for Network Basic Input Output System, an application programming interface (API) that augments the DOS BIOS by adding special functions for local-area networks (LANs). Almost all LANs for PCs are based on the NetBIOS. Some LAN manufacturers have even extended it, adding additional network capabilities.

The Avaya P330 can be configured to relay netbios UDP broadcast packets. This feature is used for applications such as WINS that use broadcast but may need to communicate with stations on other subnets or VLANs.

Configuration is performed on a per-interface basis. When a netbios broadcast packet arrives from an interface on which netbios rebroadcast is enabled, the packet is distributed to all other interfaces configured to rebroadcast netbios.

If the netbios packet is a net-directed broadcast (e.g., 149.49.255.255), the packet is relayed to all other interfaces on the list, and the IP destination of the packet is replaced by the appropriate interface broadcast address.

If the netbios broadcast packet is a limited broadcast (e.g., 255.255.255.255), it is relayed to all VLANs on which there are netbios-enabled interfaces. In that case, the destination IP address remains the limited broadcast address.

NetBIOS Re-broadcast Configuration CLI Commands

In order to...	Use the following command...
Set NetBIOS rebroadcasts mode on an interface	ip netbios-rebroadcast
Disable NetBIOS rebroadcasts mode on an interface	no ip netbios-rebroadcast

Policy Configuration

Policy Configuration Overview

The P330 supports QoS (Quality of Service) by using multiple priority levels and IEEE 802.1p priority tagging to ensure that data and voice receive the necessary levels of service.

The Avaya P330 can enforce QoS policy on routed packets and change their 802.1p priority, according to the following criteria:

- The packet protocol
- Matching the packet's source or destination IP address to the configured priority policy.
- Whether the packet source or destination TCP/UDP port number falls within a pre-defined range.

In addition, the 802.1p priority of a packet can be modified according to the DSCP value in the IP header based on the DSCP-802.1p mapping configured by the user.

The P330 supports Access Control policy. Access Control rules define how the P330 should handle routed packets. There are three possible ways to handle such packets:

- Forward the packet (Permit operation)
- Discard the packet (Deny operation)
- Discard the packet and notify the management station (Deny and Notify)

The Avaya P330 can enforce Access Control policy on each routed packet, according to the following criteria:

- Matching the packet's source or destination IP address to the configured Access Control policy.
- Determine if the packet protocol and source or destination TCP/UDP port number falls within a pre-defined range.
- Using the ACK bit of the TCP header.

The P330 uses policy lists containing both Access Control rules and QoS rules. The policy lists are ordered by rule indexing.

The Avaya P330 access control rules are set-up using the Command Line Interface and Avaya EZ2Rule central policy management application under Avaya™ MSNM (MultiService Network Manager).

Enforcement

P333R-LB can enforce policy rules on traffic addressed to its interfaces. This feature enables you to block any configuration (e.g., SNMP, TELNET, and HTTP) of the router/load balancer.

Default List Behavior

P333R-LB has a default list which is always active. The default list has one implicit rule permitting all traffic to be forwarded. The DSCP table of the default list is activated together with the default list, and includes the default DSCP-CoS mappings.

Policy and Load Balancing

P333R-LB can enforce policy rules relating to Virtual IP addresses and to the original Client IP. P333R-LB can not enforce policy rules relating to PIP addresses.

When a packet destined to a Virtual IP address is processed by P333R-LB, it checks the packet against the active policy list before any NAT manipulation is performed on the packet. Upon returning from the real servers, the packet is checked against the active policy list after NAT was performed.

P333R-LB Policy Properties

P333R-LB supports QoS and Access Control rules pertaining to source IP addresses, destination IP addresses, any IP protocol, any specific TCP/UDP port, and the ACK bit of the TCP header. It also supports QoS rules on up to 3 port-ranges of TCP/UDP.

Policy Configuration CLI Commands

In order to...	Use the following command...
Display the DSCP-802.1p mapping	show dscp
Activate a specific policy list	ip access-group
Deactivate a specific policy list	no ip access-group
Set the default action for a specific policy list	ip access-default-action
Set a name for a policy list	ip access-list-name
Set the owner for a specific policy list	ip access-list-owner

In order to...	Use the following command...
Create a specific policy rule	ip access-list
Delete a specific policy rule	no ip access-list
Check the policy for a simulated packet	ip simulate
Set the list cookie for a specific policy list	ip access-list-cookie
Copy a configured source policy list to a destination policy list	ip access-list-copy
Set the scope of a policy list	ip access-list-scope
Verify that all the rules in a priority list are valid	validate-group
Display information about the configured active access list.	show ip access-group
Display all the current policy lists	show ip-access lists
Display a complete data for a policy list (including DSCP and composite operation tables)	show ip-access-lists-details
Set the policy control source to either local or remote policy server	set qos policy-source

Policy Configuration Example

Figure 12.3 Avaya P330 Policy

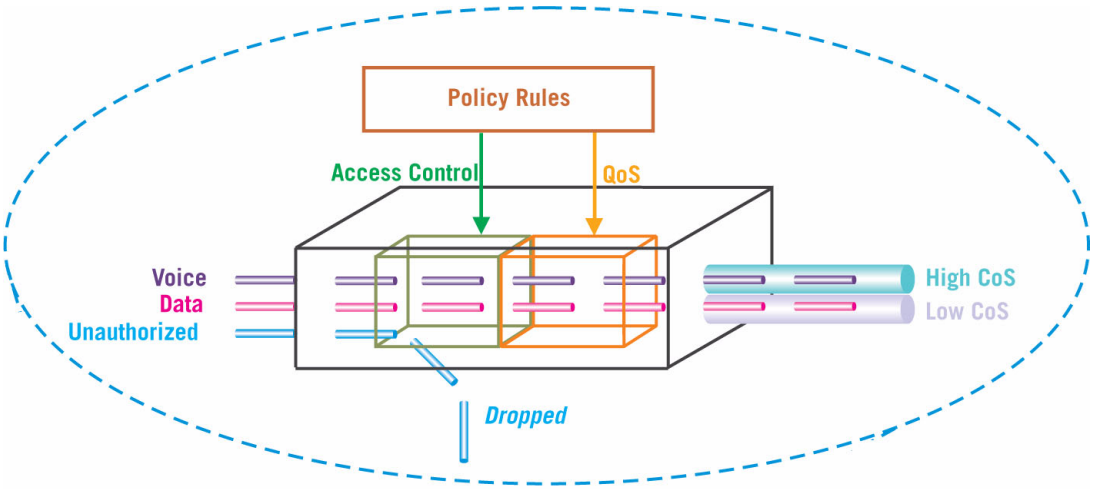

Policy Configuration Example

The following shows configuration of Access List 100:

1 Assigning priority 6 to all TCP traffic originating in network 149.49.0.0 – rule 1:

```
P330-1(super)# ip access-list 100 1 fwd6 tcp 149.49.0.0
0.0.255.255 any
done!
```

2 Assigning priority 3 to all TCP traffic going to the host 172.44.17.1 – rule 2:

```
P330-1(super)# ip access-list 100 2 fwd3 tcp any host
172.44.17.1
done!
```

3 Denying Telnet sessions originated by the host 192.168.5.33 – rule 3:

```
P330-1(super)# ip access-list 100 3 deny tcp host
192.168.5.33 any eq 23
done!
```

IP Fragmentation and Reassembly

IP Fragmentation and Reassembly Overview

The P330 supports IP Fragmentation and Reassembly. This feature allows the router to send and receive large IP packets where the underlying data link protocol constrains MTU (maximum transport unit).

IP fragmentation involves breaking a datagram into a number of pieces that can be reassembled later. The IP source, destination, identification, total length, and fragment offset fields, along with the "more" fragment and "don't" fragment flags in the IP header, are used for IP fragmentation and reassembly.

IP Fragmentation works as follows:

- IP packet is divided into fragments
- each fragment becomes its own IP packet
- each packet has same identifier, source, destination address
- fragments are usually not reassembled until final destination

IP Fragmentation/Reassembly CLI Commands

In order to...	Use the following command...
Clear the fragment database and restore its defaults	clear fragment
Set the maximum number of fragments that can comprise a single IP packet	fragment chain
Set the maximum number of fragmented IP packets, destined for the router, to reassemble at any given time	fragment size
Set the maximum number of seconds to reassemble a fragmented IP packet destined for the router.	fragment timeout
Display information regarding fragmented IP packets that are destined for the router	show fragment

Layer 3 Configuration File

The Configuration File feature allows the user to read the routing configuration parameters and save them to a file on the station. The routing configuration commands in the file are in CLI format. The user can edit the file (if required) and re-configure the router module by downloading the configuration file.

Although the file can be edited, it is recommended to keep changes to the file to a minimum. The recommended configuration method is using MSNM P330 Device Manager and/or the CLI. Changes to the configuration file should be limited to those required to customize a configuration file from one router to suit another.

Chapter 13

Layer 3 Redundancy

This chapter discusses the redundancy schemes of VRRP, SRRP, and additional redundancy schemes (Real Server Group Backup and Real Server Backup).

VRRP

VRRP is an IETF protocol designed to support redundancy of routers on the LAN, as well as load balancing of traffic. VRRP is transparent to host stations, making it an ideal choice when redundancy, load balancing and ease of configuration are all required.

The concept underlying VRRP is that a router can backup other routers, in addition to performing its primary routing functions. Redundancy is achieved by introducing the concept of a virtual router. A virtual router is a routing entity associated with multiple physical routers. The routing functions of the virtual router are performed by one of the physical routers with which it is associated. This router is known as the master router.

For each virtual router, VRRP selects a master router. If the selected master router fails, another router is selected as master router.

In VRRP, two or more physical routers can be associated with a virtual router, thus achieving the extreme reliability inherent in the P333R-LB SAFER architecture.

In a VRRP environment, host stations interact with the virtual router. They are not aware that this router is a virtual router, and they are not affected when a new router takes over the role of master router. This makes VRRP fully interoperable with every host station.

VRRP can be activated on an interface using a single command, while allowing for the necessary fine-tuning of the many VRRP parameters. For a detailed description of VRRP, refer to VRRP standards and published literature.

Introducing a Load Balancer into the network creates a single point of failure. As a result, users will most likely want a backup implementation between two P333R-LBs.

For FWLB, the `my-ip` parameter of the `hc-ip` command can also be an associated IP address of a virtual router. In addition, the `override addr owner` parameter of the `ip vrrp` command should be used.

For AR, at least two virtual routers should be configured: one for the client vlan area and the other for the server vlan area. The servers should be configured with the VRRP IP as their default gateway. In addition, the clients should be configured with the VRRP IP as their default gateway. The same physical router should be the

master of all the virtual routers, using the `priority` command if necessary.

For SLB, a modification to VRRP is necessary. In this case, you need to backup the VIP (in addition to a routing interface of the router). An interface on the VIP's subnet still has to be configured. The VIP will be configured as the Associated IP of the VRRP.

Note: When two P333R-LBs are connected in a VRRP pair with Full NAT configured on both, you need to configure different PIP address ranges for each P333R-LB.

VRRP Commands

The following table displays the VRRP Commands:

Figure 13.1 VRRP Commands

In order to...	Use the following command...
Display VRRP information	show ip vrrp
Display full VRRP-related information	show ip vrrp detail
Enable/disable VRRP routing globally	router vrrp
Create/delete a virtual router on the interface	ip vrrp
Assign/remove an IP address to the virtual router on the interface	ip vrrp address
Set the virtual router advertisement timer value (in seconds) for the virtual router ID	ip vrrp timer
Set the virtual router priority value used when selecting a master router	ip vrrp priority
Set the virtual router simple password authentication for the virtual router ID	ip vrrp auth-key
Configure the router to preempt a lower priority master for the virtual router ID	ip vrrp preempt

In order to...	Use the following command...
Set the primary address used as the source address of VRRP packets for the virtual router ID	ip vrrp primary
Accept packets addressed to the IP address(es) associated with the virtual router	ip vrrp override addr owner

Configuration Example

The following figure illustrates a VRRP configuration according to SLB.

Figure 13.2 SLB - VRRP Configuration Example

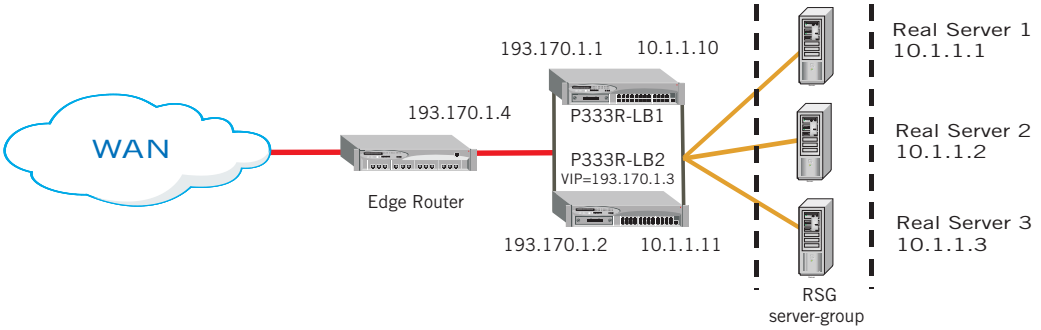

The following configuration file is a result of the above configuration:

```
--------------
P333R-LB1
--------------
hostname "MASTER"
!
router vrrp
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.10       255.255.255.0
 enable vlan commands
 ip vrrp 1
 ip vrrp 1 address 10.1.1.10
 ip vrrp 1 address 193.170.1.1
 ip vrrp 2
 ip vrrp 2 priority 254
 ip vrrp 2 address 193.170.1.3
!
```

```
ip default-gateway 193.170.1.4     1 low
!
real-slb-server 10.1.1.1
!
real-slb-server 10.1.1.2
!
real-slb-server 10.1.1.3
 !
rsg "server-group"
 type-id slb 1
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
 real-slb-server 10.1.1.3
!
virtual-server "web-farm"
 id 1
 vip 193.170.1.3
 !
 virtual-slb-service "www-service"
  id 1
  application tcp 80
  rsg "server-group"

--------------
P333R-LB2
--------------
hostname "BACKUP"
!
router vrrp
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.2     255.255.255.0
!
interface "2"
 ip vlan name "Default"
```

```
 ip address 10.1.1.11       255.255.255.0
 enable vlan commands
 ip vrrp 1
 ip vrrp 1 address 10.1.1.10
 ip vrrp 1 address 193.170.1.1
 ip vrrp 2
 ip vrrp 2 address 193.170.1.3
!
ip default-gateway 193.170.1.4     1 low
!
real-slb-server 10.1.1.1
!
real-slb-server 10.1.1.2
!
real-slb-server 10.1.1.3
 !
rsg "server-group"
 type-id slb 1
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
 real-slb-server 10.1.1.3
!
virtual-server "web-farm"
 id 1
 vip 193.170.1.3
 !
 virtual-slb-service "www-service"
  id 1
  application tcp 80
  rsg "server-group"
```

SRRP

P333R-LB SRRP redundancy capabilities provide automatic backup Layer 3 switching for IP stations. P333R-LB units can be configured to back each other up so that if one fails the other will take over its forwarding functions. The backup P333R-LB is not idle. As long as both P333R-LB units are functional, traffic is shared between them. The P333R-LB modules can be in the same P330 stack or in different, connected, P330 stacks. The P333R-LB can back up another P333R-LB unit or any other router.

A P333R-LB unit configured to back up another unit monitors the other's status by polling it at configured intervals, and automatically detects when the other router fails and when it becomes functional again. When detecting a failure, the backup P333R-LB sends a gratuitous ARP message that causes all stations to send their IP traffic to the backup P333R-LB MAC address instead of the failed unit MAC address. As long as it is an active backup resulting from the failure of the main unit, the backup P333R-LB answers ARP requests for the main unit, providing its own MAC address.

SRRP Commands

The following table displays the SRRP Commands:

Figure 13.3 SRRP Commands

In order to...	Use the following command...
Display the SRRP configuration and status	show ip srrp
Configure SRRP options, activate SRRP and enter the SRRP configuration mode	router srrp
Backup an additional interface of the main router using the SRRP application	ip srrp backup
Configure the polling interval (in seconds) used by SRRP	poll-interval
Configure the timeout (in seconds) after which SRRP declares the main router dead if it does not respond to polling	timeout

Additional Redundancy Schemes

Besides VRRP and SRRP, the P333R-LB implements two additional redundancy schemes:

- Real Server Group Backup.
- Real Server Backup.

Real Server Group Backup

You can configure a Real Server Group (RSG) to serve as a backup for a primary RSG implementing a Virtual Service. You can configure an RSG to backup one or more primary RSGs.

An RSG can backup a primary RSG, but not a backup RSG. In addition, you cannot have two RSGs backing up each other.

A backup RSG can run a different Virtual Service than the primary RSG while supplying backup to all services of the primary RSG.

The backup RSG is not used for the primary RSG's services until all the Real Servers in the primary RSG are down. When backup is implemented, the backup RSG runs the primary RSG's service in addition to its own services.

When the primary RSG has recovered, it resumes operation. Ongoing sessions continue going to the backup RSG until they are completed. In addition, in order to maintain persistency, new sessions might open to the backup RSG.

RSG Backup Configuration Example

The following figure illustrates an RSG Backup configuration.

Figure 13.4 RSG Backup Configuration Example

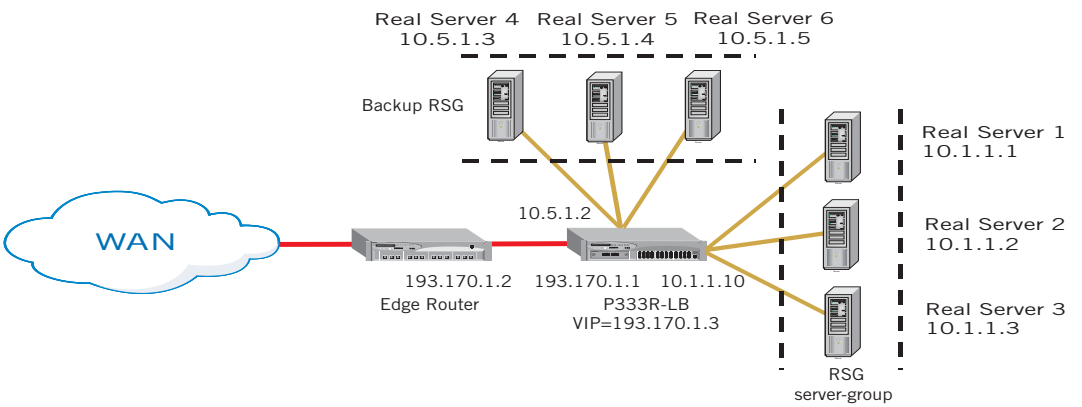

The following configuration file is a result of configuring your network as shown in Figure 13.4:

```
-------------
P333R-LB
-------------

hostname "P333R-LB"
!
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.10       255.255.255.0
!
interface "3"
 ip vlan name "Default"
 ip address 10.5.1.2        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
!
real-slb-server 10.1.1.1
!
real-slb-server 10.1.1.2
!
real-slb-server 10.1.1.3
!
real-slb-server 10.5.1.3
!
real-slb-server 10.5.1.4
!
real-slb-server 10.5.1.5
!
rsg "backup-group"
 type-id slb 1
```

```
 real-slb-server 10.5.1.3
 real-slb-server 10.5.1.4
 real-slb-server 10.5.1.5
!
rsg "server-group"
 type-id slb 2
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
 real-slb-server 10.1.1.3
!
rsg "server-group"
 backup 1
!
virtual-server "web-farm"
 id 1
 vip 193.170.1.3
 !
 virtual-slb-service "www-service"
  id 1
  application tcp 80
  rsg "server-group"
```

Real Server Backup

You can configure a Real Server to serve as a backup for one or more primary Real Servers in a Real Server Group.

As with RSG, a Real Server can backup a primary Real Server but not a backup Real Server. In addition, you cannot have two Real Servers backing up each other.

Unlike a backup RSG, a backup Real Server cannot provide another service while supplying backup to the primary Real Server. Although the backup Real Server will be completely dormant while the primary Real Server is active, its Health will be periodically checked.

A backup Real Server will not be used until the primary Real Server is down.

Note: A backup Real Server cannot be a part of an RSG.

When the primary Real Server has recovered, it will resume operation and begin to receive new sessions. Ongoing sessions will continue going to the backup Real Server until they are completed. In addition, to keep persistency, new sessions might open to the backup Real Server.

Real Server Backup Configuration Example

The following figure illustrates a Real Server Backup configuration.

Figure 13.5 Real Server Backup Configuration Example

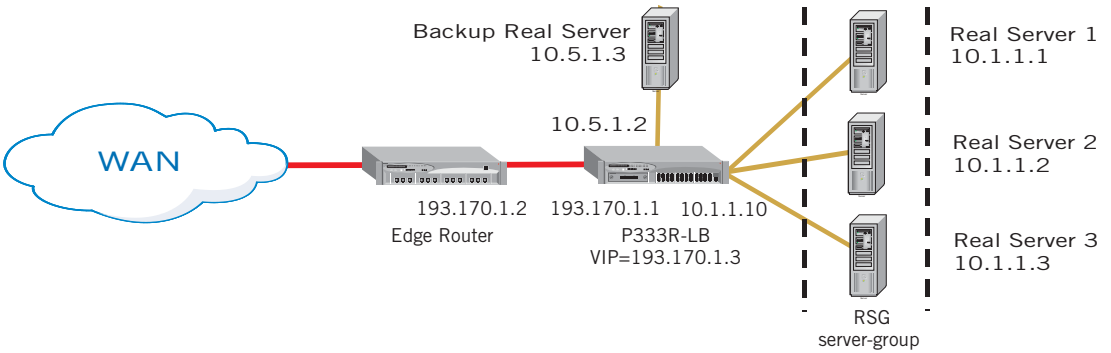

The following configuration file is a result of configuring your network as shown in Figure 13.5:

```
----------------
P333R-LB
----------------
hostname "P333R-LB"
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.10       255.255.255.0
!
interface "3"
 ip vlan name "Default"
 ip address 10.5.1.2        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
!
real-slb-server 10.5.1.3
!
real-slb-server 10.1.1.1
 backup 10.5.1.3
!
real-slb-server 10.1.1.2
 backup 10.5.1.3
!
real-slb-server 10.1.1.3
 backup 10.5.1.3

!
rsg "server-group"
 type-id slb 1
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
```

```
 real-slb-server 10.1.1.3
!
virtual-server "web-farm"
 id 1
 vip 193.170.1.3
 !
 virtual-slb-service "www-service"
  id 1
  application tcp 80
  rsg "server-group"
```

Chapter 14

Load Balancing in the P333R-LB

This chapter provides information on Load Balancing in the Avaya P333R-LB. Both Firewall Load Balancing and Server Load Balancing are treated in detail.

Firewall Load Balancing

This section provides information on Firewall Load Balancing (FWLB) in general, as well as on specific types of firewalls supported by the P333R-LB, and configuration examples.

FWLB intercepts all traffic between protected and unprotected zones, and dynamically distributes the load among the available firewalls, based on the FWLB configuration.

In terms of P333R-LB configuration, firewalls are referred to as Real Servers, the group of firewalls is a Real Server Group, and the firewall group is associated with a Virtual Firewall Service, which is a routing or bridging firewall.

Benefits

FWLB allows multiple firewalls to operate in parallel, giving you the ability to:

- Scale firewall performance.
- Eliminate the firewall as the single point of failure.

How It Works

The P333R-LB:

- Balances traffic across two or more firewalls (up to1024) in your network, allowing your firewalls to work in parallel.
- Maintains state information about the traffic flowing through it and ensures that all traffic between specific IP source and destination pairs flows through the same firewall.
- Performs health checks on all paths through the firewalls. If a path is malfunctioning, P333R-LB diverts traffic away from that path, maintaining connectivity across the firewalls.

Often, a pair of P333R-LBs are needed to support FWLB. One device is deployed on the LAN side (internal) of the firewalls and another on the WAN side (external). Another P333R-LB is required to balance a Demilitarized Zone (DMZ) on the DMZ side of the network, if it exists, to allow remote access. Additional P333R-LBs can be added to provide redundancy, eliminating any P333R-LB or path as the single point

of failure.

P333R-LB supports both bridging and two types of routing firewalls: transparent and non-transparent.

Transparent Routing Firewall Load Balancing

This section explains how the P333R-LB supports Transparent Routing FWLB, and includes configuration examples.

Implementation

For Transparent Routing FWLB, the load balancer receives a packet, makes a load balancing decision, and forwards the packet to a firewall. The firewall does not perform Network Address Translation (NAT) on the packets; the source and destination IP addresses are not changed.

Two P333R-LBs are required for Transparent Routing FWLB, one on each side of the firewalls. One device intercepts traffic between the protected zone and the firewall, and the second device intercepts traffic between the unprotected zone and the firewall.

Transparent Routing firewalls act as "next hop" devices from the perspective of the P333R-LB. After one of the firewalls in a group is selected, normal routing to that firewall takes place.

The P333R-LB performs an “intelligent routing” decision, based on the Load Balancing criteria, and replaces the MAC address in the packets by the MAC address of the selected firewall.

P333R-LB enables you to route packets destined to a Demilitarized Zone (DMZ). A DMZ is a portion of the client’s network, apart from the client’s LAN, where remote access is allowed. After creating a DMZ, a third load balancer is installed, which routes packets to the DMZ.

Configuring Firewall Load Balancing in the P333R-LB

The following figure illustrates a simple transparent routing FWLB configuration.

Figure 14.1 Transparent Routing FWLB Sample Configuration

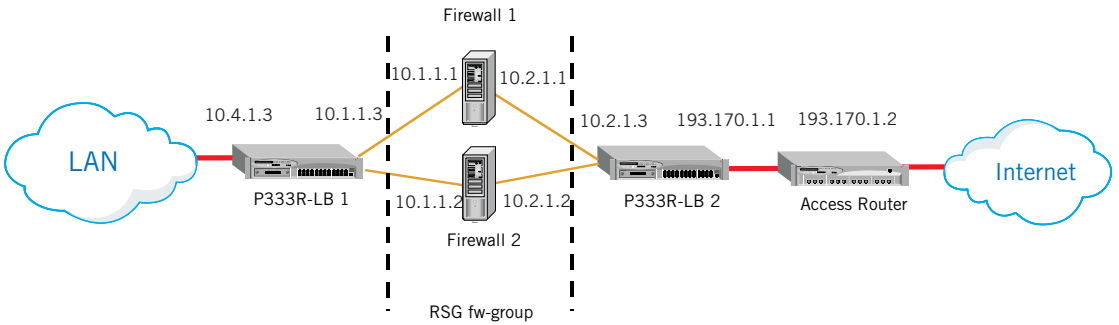

Note:
1. When configuring routing firewalls as Real Servers, you must give an ID to each Real Server. This ID must match the ID given to the same firewall on the second load balancer.
2. The P333R-LB performs load balancing on traffic that arrives to its routing interfaces. Therefore, IP routes in the network must be configured to pass through the P333R-LB.

To configure your network as in Figure 14.1, the following should be done:

- The LAN routers (or hosts) should be configured with 10.4.1.3 as the next hop toward the WAN (the default gateway in many cases).
- The access router should be configured with 193.170.1.1 as the next hop toward the LAN.
- The firewalls should be configured with 10.1.1.3 as the next hop towards the LAN, and 10.2.1.3 as the next hop toward the WAN.
- The firewalls must be configured to allow ICMP Echo to pass between the two load balancers (10.1.1.3 and 10.2.1.3) for health-check purposes.

In order to configure P333R-LB1 according to Figure 14.1, perform the following commands:

```
P330-1(configure)# session router
Router-1(configure)# hostname "P333R-LB 1"
P333R-LB 1-1(configure)# interface 1
Done!
P333R-LB 1-1(config-if:1)# ip address 10.4.1.3 255.255.255.0
Done!
P333R-LB 1-1(config-if:1)# exit
P333R-LB 1-1(configure)# interface 2
Done!
P333R-LB 1-1(config-if:2)# ip address 10.1.1.3 255.255.255.0
Done!
P333R-LB 1-1(config-if:2)# exit
P333R-LB 1-1(configure)# real-routing-fw 10.1.1.1
Done!
P333R-LB 1-1(config-rsrvr:10.1.1.1)# id 1
Done!
P333R-LB 1-1(config-rsrvr:10.1.1.1)# exit
P333R-LB 1-1(configure)# real-routing-fw 10.1.1.2
Done!
P333R-LB 1-1(config-rsrvr:10.1.1.2)# id 2
Done!
P333R-LB 1-1(config-rsrvr:10.1.1.2)# exit
P333R-LB 1-1(configure)# rsg fw-group
Done!
P333R-LB 1-1(config-rsg:fw-group)# type-id routing-fw 1
Done!
P333R-LB 1-1(config-rsg:fw-group)# real-routing-fw 10.1.1.1
Done!
P333R-LB 1-1(config-rsg:fw-group)# real-routing-fw 10.1.1.2
Done!
P333R-LB 1-1(config-rsg:fw-group)# exit
P333R-LB 1-1(configure)# virtual-fw-service external
Done!
```

```
P333R-LB 1-1(config-v-fw-srvc:external)# id 1
Done!
P333R-LB 1-1(config-v-fw-srvc:external)# hc-ip 10.1.1.3
10.2.1.3
Done!
P333R-LB 1-1(config-v-fw-srvc:external)# ip route 0.0.0.0
0.0.0.0
Done!
P333R-LB 1-1(config-v-fw-srvc:external)# rsg fw-group
Done!
P333R-LB 1-1(config-v-fw-srvc:external)# exit
```

In order to configure P333R-LB2 according to Figure 14.1, perform the following commands:

```
P330-1(configure)# session router
Router-1(configure)# hostname "P333R-LB 2"
P333R-LB 2-1(configure)# interface 1
Done!
P333R-LB 2-1(config-if:1)# ip address 193.170.1.1
255.255.255.0
Done!
P333R-LB 2-1(config-if:1)# exit
P333R-LB 2-1(configure)# interface 2
Done!
P333R-LB 2-1(config-if:2)# ip address 10.2.1.3 255.255.255.0
Done!
P333R-LB 2-1(config-if:2)# exit
P333R-LB 2-1(configure)# ip default-gateway 193.170.1.2
Done!
P333R-LB 2-1(configure)# real-routing-fw 10.2.1.1
Done!
P333R-LB 2-1(config-rsrvr:10.2.1.1)# id 1
Done!
P333R-LB 2-1(config-rsrvr:10.2.1.1)# exit
P333R-LB 2-1(configure)# real-routing-fw 10.2.1.2
```

```
Done!

P333R-LB 2-1(config-rsrvr:10.2.1.2)# id 2
Done!
P333R-LB 2-1(config-rsrvr:10.2.1.2)# exit
P333R-LB 2-1(configure)# rsg fw-group
Done!
P333R-LB 2-1(config-rsg:fw-group)# type-id routing-fw 1
Done!
P333R-LB 2-1(config-rsg:fw-group)# real-routing-fw 10.2.1.1
Done!
P333R-LB 2-1(config-rsg:fw-group)# real-routing-fw 10.2.1.2
Done!
P333R-LB 2-1(config-rsg:fw-group)# exit
P333R-LB 2-1(configure)# virtual-fw-service internal
Done!
P333R-LB 2-1(config-v-fw-srvc:internal)# id 1
Done!
P333R-LB 2-1(config-v-fw-srvc:internal)# hc-ip 10.2.1.3
10.1.1.3
Done!
P333R-LB 2-1(config-v-fw-srvc:internal)# ip route 10.4.1.0
255.255.255.0
Done!
P333R-LB 2-1(config-v-fw-srvc:internal)# rsg fw-group
Done!
P333R-LB 2-1(config-v-fw-srvc:internal)# exit
```

The following configuration file is a result of the above configuration:

```
---------------
P333R-LB 1
----------------
hostname "P333R-LB 1"
!
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.3        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 10.4.1.3        255.255.255.0
!
real-routing-fw 10.1.1.1
 id 1
!
real-routing-fw 10.1.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.1.1.1
 real-routing-fw 10.1.1.2
!
virtual-fw-service "external"
 id 1
 hc-ip 10.1.1.3 10.2.1.3
 ip route 0.0.0.0 0.0.0.0
 rsg "fw-group"
```

```
----------------
P333R-LB 2
----------------
hostname "P333R-LB 2"
!
interface "2"
 ip vlan name "Default"
 ip address 10.2.1.3        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
!
real-routing-fw 10.2.1.1
 id 1
!
real-routing-fw 10.2.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.2.1.1
 real-routing-fw 10.2.1.2
!
virtual-fw-service "internal"
 id 1
 hc-ip 10.2.1.3 10.1.1.3
 ip route 10.4.1.0 255.255.255.0
 rsg "fw-group"
```

Demilitarized Zone (DMZ) Configuration Example

The following figure illustrates Transparent FWLB with DMZ configuration.

Figure 14.2 Transparent Routing FWLB Sample DMZ Configuration

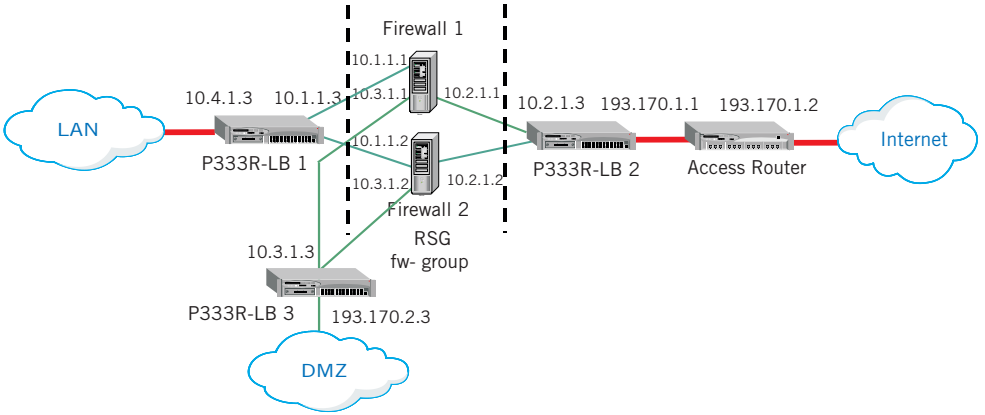

Note:

1. When configuring routing firewalls as Real Servers, you must give an ID to each Real Server. This ID must match the ID given to the same firewall on the second load balancer.
2. The P333R-LB performs load balancing on traffic that arrives to its routing interfaces. Therefore, IP routes in the network must be configured to pass through the P333R-LB.

To configure your network as in Figure 14.2, the following should be done:

- The LAN routers (or hosts) should be configured with 10.4.1.3 as the next hop toward the WAN (the default gateway in many cases).
- The access router should be configured with 193.170.1.1 as the next hop toward the LAN.
- The firewalls should be configured with 10.1.1.3 as the next hop towards the LAN, and 10.2.1.3 as the next hop toward the WAN (internet).
- The firewalls must be configured to allow ICMP Ping to pass between the two load balancers (10.1.1.3 and 10.2.1.3) for health-check purposes.
- Each load balancer must be configured to two virtual firewall services. In Figure 14.2, P333R-LB1 should be assigned to the WAN and DMZ, P333R-LB2 to the LAN and DMZ, and P333R-LB3 to the LAN and WAN.

The following configuration file is a result of configuring your network as shown in Figure 14.2:

```
--------------------
P333R-LB 1
--------------------
hostname "P333R-LB 1"
!
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.3        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 10.4.1.3        255.255.255.0
!
real-routing-fw 10.1.1.1
 id 1
!
real-routing-fw 10.1.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.1.1.1
 real-routing-fw 10.1.1.2
!
virtual-fw-service "external"
 id 1
 hc-ip 10.1.1.3 10.2.1.3
 ip route 0.0.0.0 0.0.0.0
 rsg "fw-group"
!
virtual-fw-service "dmz"
 id 2
 hc-ip 10.1.1.3 10.3.1.3
 ip route 193.170.2.0 255.255.255.0
```

```
 rsg "fw-group"

-------------------
P333R-LB 2
-------------------
hostname "P333R-LB 2"
!
interface "1"
 ip vlan name "Default"
 ip address 10.2.1.3        255.255.255.0
!
interface "2"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
!
real-routing-fw 10.2.1.1
 id 1
!
real-routing-fw 10.2.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.2.1.1
 real-routing-fw 10.2.1.2
!
virtual-fw-service "internal"
 id 1
 hc-ip 10.2.1.3 10.1.1.3
 ip route 10.4.1.0 255.255.255.0
 rsg "fw-group"
!
virtual-fw-service "dmz"
 id 2
```

```
 hc-ip 10.2.1.3 10.3.1.3
 ip route 193.170.2.0 255.255.255.0
 rsg "fw-group"

-------------------
P333R-LB 3
-------------------

hostname "P333R-LB 3"
!
interface "1"
 ip vlan name "Default"
 ip address 10.3.1.3        255.255.255.0
!
interface "2"
 ip vlan name "Default"
 ip address 193.170.2.3     255.255.255.0
!
real-routing-fw 10.3.1.1
 id 1
!
real-routing-fw 10.3.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.3.1.1
 real-routing-fw 10.3.1.2
!
virtual-fw-service "internal"
 id 1
 hc-ip 10.3.1.3 10.1.1.3
 ip route 10.4.1.0 255.255.255.0
 rsg "fw-group"
!
virtual-fw-service "external"
```

```
id 2
hc-ip 10.3.1.3 10.2.1.3
ip route 0.0.0.0 0.0.0.0
rsg "fw-group"
```

Load Balancing Metrics for Transparent Routing FWLB

There are different methods, or metrics, that the P333R-LB can use to distribute traffic among multiple firewalls. These metrics tell the P333R-LB which server should receive the next session.

Transparent Routing FWLB uses the following metrics:

- Hash.
- MinMiss Hash.

Each load balancing metric can be performed on one of the following metric parameters:

- SRC IP (source IP address).
- DST IP (destination IP address).
- SRC and DST IP (a combination of source and destination IP addresses).

Hash

Hash is the default metric for FWLB, and the default metric parameter for Transparent FWLB Hash is: Source IP, Destination IP.

Using the Hash metric, sessions are distributed through firewalls using a predefined mathematical hash function. This hash function is created by a hash value, performed on a specific parameter. The source and destination IP addresses are used as the hash function input.

P333R-LB creates a list of all the currently available firewalls. The result of the hash function is used to select a firewall from the list. Specific source and destination IP addresses always produce the same hash function, providing natural persistency.

If a firewall is added or removed to the group, the persistency will be broken. This occurs, since the order of the firewalls in the list changes, but the hash still points to the same list entries. The following figure illustrates how a loss of persistency occurs when a firewall becomes non-operational.

Figure 14.3 Hash Metric - Loss of Persistency

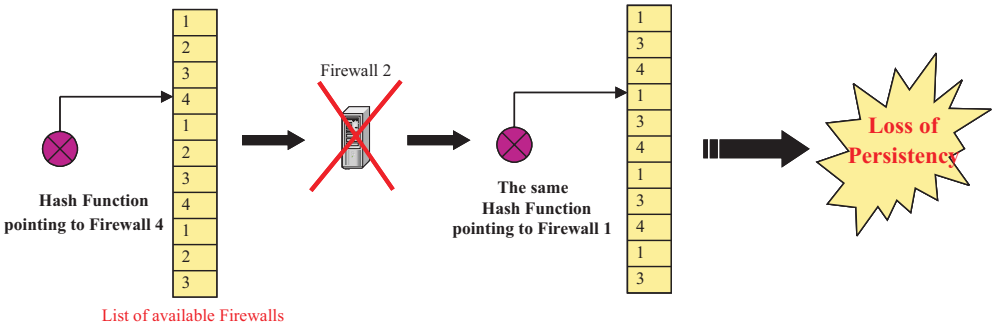

When Firewall 2 is removed from the group, the list of available firewalls is readjusted, causing a lack of persistency. However, if Firewall 2 becomes operational again, the list of available firewalls is restored to its original order, and persistency is recovered, making Hash a predictable metric.

MinMiss Hash

The metric parameter for MinMiss Hash is: Source IP, Destination IP.

MinMiss Hash distributes sessions to firewalls in the same way as the Hash metric. However, MinMiss Hash sustains persistency even when a server is removed from the firewall group. When a firewall fails or is removed, the positions of the other firewalls in the list do not change. Instead, the remaining firewalls are redistributed to the list entries freed by the failing firewall. The following figure illustrates how persistency is maintained even though a firewall becomes non-operational.

Figure 14.4 MinMiss Hash Metric - Persistency Sustained

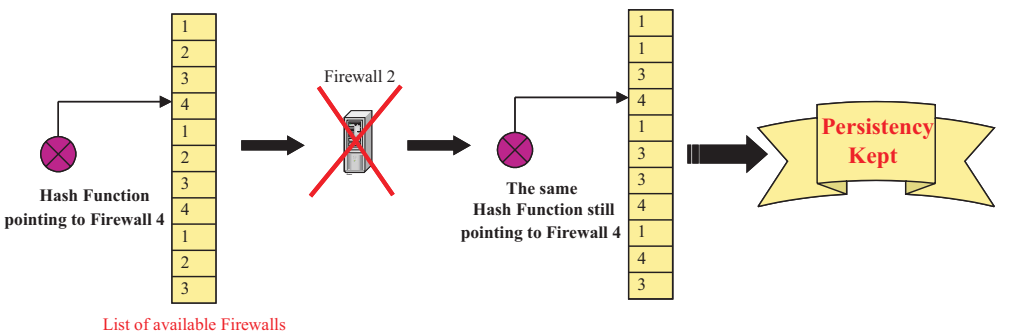

When Firewall 2 is removed from the group, the list of available firewalls is not readjusted. Only the list entries that are now empty are replaced with other available firewalls. Therefore, persistency is sustained for all available firewalls. If Firewall 2 becomes operational again, the list of available firewalls is recalculated so that the smallest number of firewalls is affected. However, the list of firewalls is not

restored to its original configuration. As a result, persistency is only partially recovered, making MinMiss Hash a non-predictable metric.

Selecting a Load Balancing Metric

The different metrics available allow you to select the metric that best matches your network topology.

The Hash metric produces predictable forwarding decisions. Therefore, this is the only metric appropriate for configurations that involve load balancers on both sides of the firewalls (for example, Transparent Routing FWLB, and Bridging FWLB). Furthermore, loss of persistency may not be an issue where a state-synchronizing firewall cluster is deployed.

For Transparent Routing FWLB, to ensure mapping of the same session to the same firewall by both load balancers, the default hash key is a combination of the source and destination IP addresses.

The MinMiss Hash metric maintains persistency better than the Hash metric. When there is not need to maintain consistency between load balancers (for example, Non Transparent FWLB), this metric should be used.

Weighted Firewalls

You can assign weights to firewalls to enable faster firewalls to receive a larger share of sessions. This minimizes overloading and maximizes functionality.

If you assign a weight to a firewall, the sessions are distributed to the firewalls in the same metric chosen (Hash or MinMiss Hash). However, weighted firewalls are assigned a larger share of sessions. For example, if you assign a weight of 10 to one firewall (the default value), and assign a weight of 20 to a second firewall, the second firewall receives 2 sessions for each session directed to the first firewall.

Health Check

The P333R-LB constantly health checks the firewall paths to ensure that each firewall is accessible and operational. A firewall that fails the health check is automatically removed from the load balancer's internal list of currently available firewalls.

The P333R-LB uses the ICMP Echo health-check method. Each load balancer periodically pings the Real Server and checks if an answer was received.

For FWLB, the health check must is performed beyond the firewalls as well in order to check the entire data path. In order to insure that the health check packets traverse the same firewall in both directions, the Health Check IP addresses (the packet's source and destination IP addresses) are the interfaces of the load balancer on each side of the firewall. For each load balancing interface, the local and remote addresses should be configured. The load balancers on both sides of the firewalls must be configured symmetrically. For information on configuration, see the

configuration example.

Persistency

Firewalls perform a Stateful Inspection on every session passing through them and drop a session if not all of its traffic passes through the same firewall. Therefore, when load-balancing between different firewalls, it is imperative that all traffic belonging to a given session traverses the same firewall.

The P333R-LB achieves this goal by implementing a sophisticated persistency mechanism, based on packet characteristics inspection. A symmetric hash function in each module is calculated based on the source and destination IP addresses. The P333R-LB assures that packets with the same characteristics traverse the same firewall in both directions throughout the session.

In the case where there are two P333R-LBs (one on each side of the firewalls), persistency is ensured only if each P333R-LB is configured so that they are compatable with each other. If they are not, and there is a change in the network that affects internal device decisions (for example, adding or removing a Real Server), persistency, or even the network connection, could be lost.

Non-Transparent Routing Firewall Load Balancing

This section explains how the P333R-LB supports non-Transparent Routing firewalls, and includes configuration examples as well.

Implementation

Non-Transparent Routing firewalls are firewalls that support dynamic NAT (Network Address Translation).

For non-Transparent FWLB, the load balancer receives an outgoing packet, makes a load balancing decision, and forwards the packet to a firewall. The firewall keeps a bank of IP addresses and replaces the source IP of the incoming packet (from the LAN) with a unique, yet arbitrary IP address from this bank. The firewall then forwards the packet to an edge router which routes it to the correct destination on the WAN.

For incoming packets, the unique NAT address is used as a destination IP to access the same firewall. The firewall performs reverse NAT by replacing the NAT destination address with the actual destination address (the client IP address), and then forwards the packet to the load balancer which routes the packet to its destination. No Load Balancing is performed on incoming packets.

For non-Transparent Routing FWLB, only one Load Balancing device is required. The device is positioned on the LAN (internal) side of the firewalls. Since the firewalls perform NAT, a Load Balancing device is not needed between the WAN and the firewalls.

As well, non-Transparent Routing FWLB can be configured using static NAT. In

this scenario, the firewalls are configured to perform NAT for some hosts, and not for other hosts. Alternatively, the firewalls may be configured to assign a specific NAT address to a specific host. In this case, two load balancers are required, one on each side of the firewalls.

NAT Configuration Example

The following figure illustrates non-Transparent Routing FWLB using NAT.

Figure 14.5 Non-Transparent Routing FWLB Sample NAT Configuration

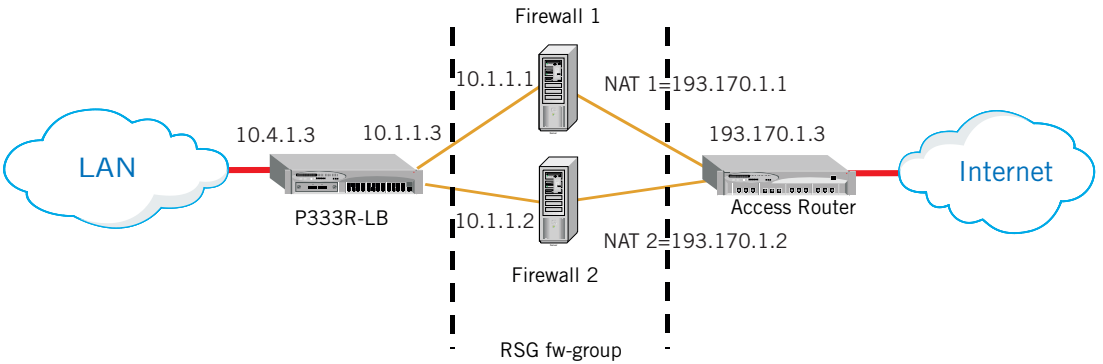

Note: The P333R-LB performs load balancing on traffic that arrives to its routing interfaces. Therefore, IP routes in the network must be configured to pass through the P333R-LB.

To configure your network as in Figure 14.5, the following should be done:

- The LAN routers (or hosts) should be configured with 10.4.1.3 as the next hop toward the WAN (the default gateway in many cases).
- The access router should be configured with 193.170.1.3 as the next hop toward the LAN.
- The firewalls should be configured with 10.1.1.3 as the next hop towards the LAN and 193.170.1.3 toward the WAN.

The following configuration file is a result of configuring your network as shown in Figure 14.5:

```
-----------------
P333R-LB
-----------------
hostname "P333R-LB"
!
```

```
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.3        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 10.4.1.3        255.255.255.0
!
real-routing-fw 10.1.1.1
 id 1
!
real-routing-fw 10.1.1.2
 id 2
!
rsg "fw-group"
 type-id routing-fw 1
 real-routing-fw 10.1.1.1
 real-routing-fw 10.1.1.2
!
virtual-fw-service "external"
 id 1
 hc-ip 10.1.1.3 193.170.1.3
 ip route 0.0.0.0 0.0.0.0
 rsg "fw-group"
```

Static NAT Configuration Example

The following figure illustrates non-Transparent Routing FWLB using static NAT.

Figure 14.6 Non-Transparent Routing FWLB Sample Static NAT Configuration

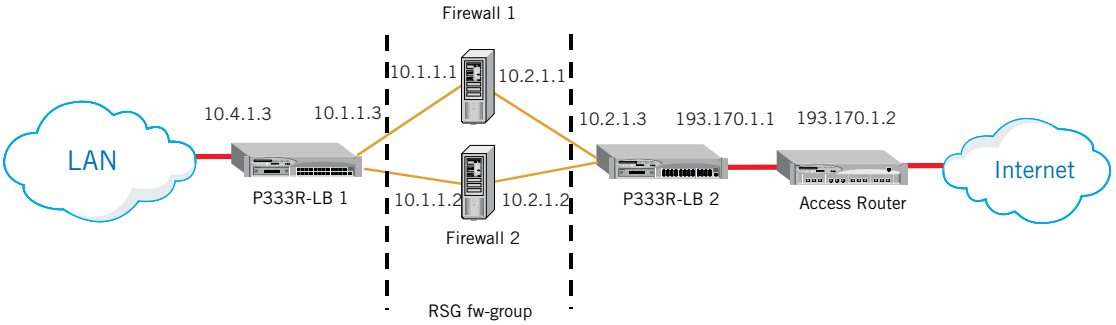

Note: The P333R-LB performs load balancing on traffic that arrives to its routing interfaces. Therefore, IP routes in the network must be configured to pass through the P333R-LB.

To configure your network as in Figure 14.6, the following should be done:

- The LAN routers (or hosts) should be configured with 10.4.1.3 as the next hop toward the WAN (the default gateway in many cases).
- The access router should be configured with 193.170.1.1 as the next hop toward the LAN.
- The firewalls should be configured with 10.1.1.3 as the next hop towards the LAN, and 10.2.1.3 as the next hop toward the WAN.
- The firewalls must be configured to allow ICMP Echo to pass between the two load balancers (10.1.1.3 and 10.2.1.3) for health-check purposes.
- Configure P333R-LB1 to use the Hash metric, and set the Hash parameter to "destination".
- Configure P333R-LB2 to use the Hash metric, and set the Hash parameter to "source".

As both load balancers compute the hash function based on IP addresses of hosts across the WAN, session persistency is maintained.

Load Balancing Metrics for Non-Transparent Routing FWLB

Non-Transparent Routing FWLB uses the following metrics:

- Hash.
- MinMiss Hash.

Hash is the default metric for FWLB, and the metric parameter is: Source IP, Destination IP. For full information about Hash and MinMiss Hash, see Load Balancing Metrics for Transparent Routing FWLB on page 13. For full information on selecting a load balancing metric, see Selecting a Load Balancing Metric on page

15.

Health Check

In order for the P333R-LB to perform a health check, an IP address beyond the firewall should be configured as a health check address. The health check session returns through the same firewall according to the NAT address it was given. For information on configuration, see the configuration example.

Persistency

As in the case of Transparent Routing Firewalls, Non-Transparent Routing Firewalls perform Statefull Inspection on all packets going through them, checking that all the packets of a given session traverse the same firewall in both directions.

In transparent FWLB, persistency is ensured by the Load Balancing device. In non-transparent FWLB, the firewalls ensure persistency through NAT, and there is no need for the Load Balancing device to intervene.

Bridging Firewall Load Balancing

This section explains how the P333R-LB supports Bridging FWLB, and includes a configuration example.

Implementation

Bridging firewalls are firewalls that do not perform forwarding at the IP layer, but rather appear as transparent bridges. Bridging firewalls are transparent to devices inside and outside the secured network.

For bridging FWLB, the P333R-LBs have to be positioned on both sides of the firewalls. The Bridging Firewall does not have an IP address or a MAC address to which traffic is directed, therefore a Bridging Firewall has to physically appear on the traffic path. Each P333R-LB load balances between IP interfaces of the peer P333R-LB behind the firewall. In order for this to work, the P333R-LB has to have a different VLAN and subnet for each firewall, and the physical ports connected to the firewalls have to be on different VLANs as well. In addition, for each VLAN, both load balancers must be in the same subnet. The Real Servers are IP interfaces of the load balancer on the other side of the firewalls, not of the firewalls themselves.

Configuration Example

The following figure illustrates Bridging FWLB.

Figure 14.7 Bridging FWLB Configuration Example

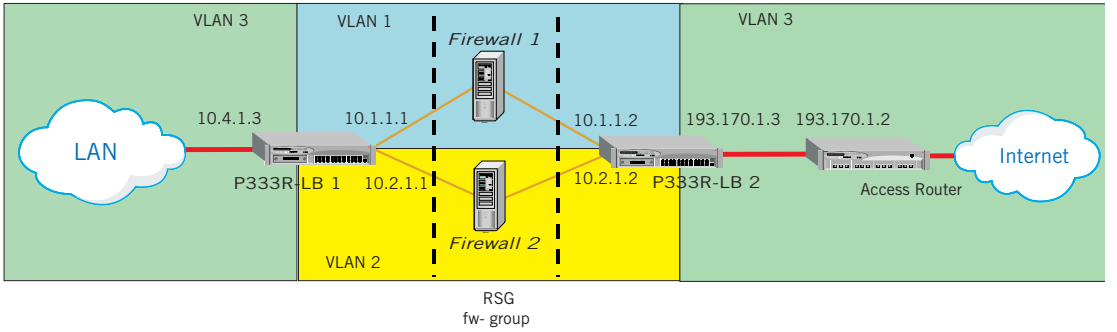

Note:
The P333R-LB performs load balancing on traffic that arrives to its routing interfaces. Therefore, IP routes in the network must be configured to pass through the P333R-LB.

To configure your network as in Figure 14.7, the following should be done:

- The LAN routers (or hosts) should be configured with 10.4.1.3 as the next hop toward the WAN (the default gateway in many cases).
- The access router should be configured with 193.170.1.3 as the next hop toward the LAN.
- The firewalls must be configured to allow ICMP Echo to pass between the two load balancers (10.1.1.1 and 10.1.1.2) for health-check purposes.

In addition, Bridging firewalls are transparent in terms of Layer 3. Therefore, to configure paths through the firewalls:

- P333R-LB1 must be configured to view the IP interfaces of P333R-LB2 (10.1.1.2, 10.2.1.2) as its "Real Servers".
- P333R-LB2 must be configured to view the IP interfaces of P333R-LB1 (10.1.1.1, 10.2.1.1) as its "Real Servers".
- Firewall 1 and Firewall 2 must be connected to different ports, and those ports must be set to different VLANs.

The following configuration file is a result of configuring your network as shown in Figure 14.7:

```
----------------
P333R-LB 1
----------------
set vlan       1 name "v1"
set vlan       2 name "v2"
set vlan       3 name "v3"
!
hostname "P333R-LB 1"
!
interface "1"
 ip vlan name "v1"
 ip address 10.1.1.1        255.255.255.0
!
interface "2"
 ip vlan name "v2"
 ip address 10.2.1.1        255.255.255.0
!
interface "3"
 ip vlan name "v3"
 ip address 10.4.1.3        255.255.255.0
!
```

```
real-bridging-fw 10.1.1.2
!
real-bridging-fw 10.2.1.2
!
rsg "fw-group"
 type-id bridging-fw 1
 real-bridging-fw 10.1.1.2
 real-bridging-fw 10.2.1.2
!
virtual-fw-service "bridging-external"
 id 1 bridging-fw
 ip route 0.0.0.0 0.0.0.0
 rsg "fw-group"

----------------
P333R-LB 2
----------------
set vlan       1 name "v1"
set vlan       2 name "v2"
set vlan       3 name "v3"
!
hostname "P333R-LB 2"
!
interface "1"
 ip vlan name "v1"
 ip address 10.1.1.2        255.255.255.0
!
interface "2"
 ip vlan name "v2"
 ip address 10.2.1.2        255.255.255.0
!
interface "3"
 ip vlan name "v3"
 ip address 193.170.1.3     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
```

```
!
real-bridging-fw 10.1.1.1
 !
real-bridging-fw 10.2.1.1
!
rsg "fw-group"
 type-id bridging-fw 1
 real-bridging-fw 10.1.1.1
 real-bridging-fw 10.2.1.1
!
virtual-fw-service "bridging-internal"
 id 1 bridging-fw
 ip route 10.4.1.0 255.255.255.0
 rsg "fw-group"
```

Load Balancing Metrics for Bridging FWLB

Bridging FWLB uses the following metrics:

- Hash.
- MinMiss Hash.

Hash is the default metric for FWLB, and the metric parameter is: Source IP, Destination IP. For full information about Hash and MinMiss Hash, see Load Balancing Metrics for Transparent Routing FWLB on page 13. For full information on selecting a load balancing metric, see Selecting a Load Balancing Metric on page 15.

Health Check

As with Transparent Routing FWLB, the interfaces on both sides of the firewall are periodically pinged and checked if an answer was received. Since each firewall is configured on a different VLAN, the ping will always remain on the same firewall.

Persistency

Each P333R-LB interface and the firewall connected to it reside in a separate VLAN. This ensures persistency since all the traffic through a particular firewall is contained in the firewall's VLAN.

Firewall Load Balancing Management Security

In some scenarios, the P333R-LB may be placed outside of the protection of the firewalls and be exposed to intrusion attempts through its L2 and L3 management interfaces. This is especially true for Transparent and Bridging Load Balancing, where the P333R-LB is placed outside the internal network. The intrusion attempts might be either via Telnet (CLI) or SNMP/HTTP (Embedded Web manager).

The user can prevent attacks by implementing the following:

- Change the L2 IP address of the stack agent to an IP address on a VLAN and subnet not accessible to the Access Router that connects the device to the outside world.
- Configure Access Rules on the IP interfaces (L3) of the P333R-LB that is exposed to the outside world, that will block Telnet, SNMP and HTTP traffic sent to those interfaces as the final destination. This model can only be managed from the LAN.

Server Load Balancing

This section provides information on Server Load Balancing (SLB) implementation of the P333R-LB, and includes configuration examples.

SLB replaces a single physical server with a group of servers that appear to the clients as a single Virtual Server.

Benefits

SLB improves your network performance by:

- Minimizing server response time.
- Maximizing server availability.
- Increasing reliability - If any server fails, the remaining servers continue to provide services seamlessly.
- Increasing scalability - Server configuration (removal/addition) can be performed without disrupting the network.

How it Works

The P333R-LB balances the traffic among several servers which all have access to identical applications and data. This involves intercepting all traffic between clients and load-balanced servers and dynamically distributing the load according to configured schemes called metrics. For more information on metrics, refer to Load Balancing Metrics on page 37.

To intercept traffic to the servers, the P333R-LB presents itself to the clients as a Virtual Server with a Virtual IP address (VIP). Client traffic travels to the P333R-LB acting as a Virtual Server. The P333R-LB redirects this traffic using Network Address Translation (NAT) to the available Real (balanced) Servers. Traffic from the Real Servers back to the clients may return via the P333R-LB, depending on the load balancing redirection method.

To maintain awareness of any changes in the Real Servers' availability, the P333R-LB implements a Health Check mechanism to monitor the status of the Real Servers. For more information, see Health Check on page 40.

P333R-LB enables the following methods of NAT, described in the following sections:

- **Half NAT** - P333R-LB translates the VIP to the IP address of the Real Servers.
- **Full NAT** - P333R-LB translates client's IP addresses to Proxy IP addresses, as well as the VIP to the IP address of the Real Servers.

Half NAT Load Balancing

In Half NAT based load balancing, the P333R-LB replaces the Virtual IP address of the P333R-LB with the real IP address of the Real Server. Responses made by the Real Server must pass through the P333R-LB, where the Real Server IP is replaced

with the Virtual Server IP before forwarding them to the clients.

Note: In complex routing topologies where alternative paths between the Real Servers and clients might exist, the returning packets could reach the client via a path external to the P333R-LB. These packets would be labelled with the real IP of the Real Server, and not the VIP that the client recognizes, causing the session to fail. In a situation where this is possible, use Full NAT to prevent this from occurring. For more information, refer to Full NAT Load Balancing on page 30.

Configuring Server Load Balancing in the P333R-LB

Half NAT Based Configuration

Note: The following example uses sample IP addresses for illustration only.

Figure 14.8 illustrates a Half NAT based SLB configuration.

Figure 14.8 Half NAT Based SLB Configuration Example

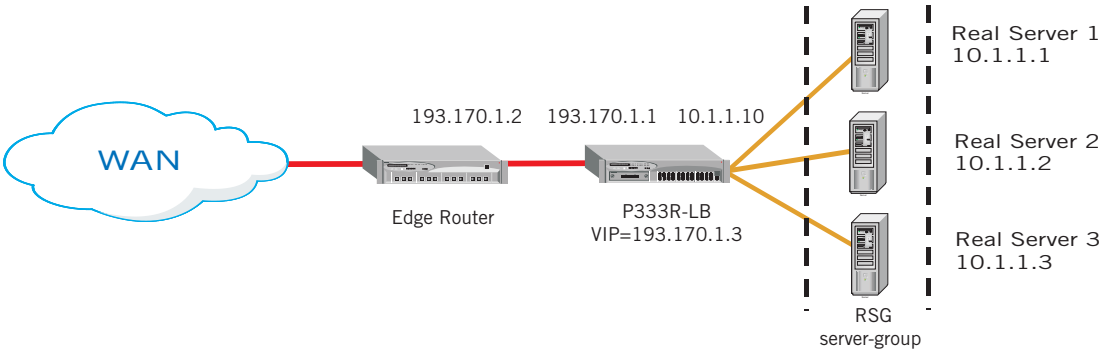

Note: The Real Servers must be configured with the P333R-LB as their default gateway (for example in Figure 14.8, 10.1.1.10).

In order to configure the load balancer according to Figure 14.8, perform the following commands:

```
P330-1(configure)# session router
Router-1(configure)# hostname "P333R-LB"
P333R-LB-1(configure)# interface 1
Done!
P333R-LB-1(config-if:1)# ip address 193.170.1.1 255.255.255.0
Done!
P333R-LB-1(config-if:1)# exit
P333R-LB-1(configure)# interface 2
Done!
P333R-LB-1(config-if:2)# ip address 10.1.1.10 255.255.255.0
Done!
P333R-LB-1(config-if:2)# exit
P333R-LB-1(configure)# ip default-gateway 193.170.1.2
Done!
P333R-LB-1(configure)# real-slb-server 10.1.1.1
Done!
P333R-LB-1(config-rsrvr:10.1.1.1)# exit
P333R-LB-1(configure)# real-slb-server 10.1.1.2
Done!
P333R-LB-1(config-rsrvr:10.1.1.2)# exit
P333R-LB-1(configure)# real-slb-server 10.1.1.3
Done!
P333R-LB-1(config-rsrvr:10.1.1.3)# exit
P333R-LB-1(configure)# rsg service-group
Done!
P333R-LB-1(config-rsg:service-group)# type-id slb 1
Done!
P333R-LB-1(config-rsg:service-group)# real-slb-server
10.1.1.1
Done!
P333R-LB-1(config-rsg:service-group)# real-slb-server
10.1.1.2
Done!
P333R-LB-1(config-rsg:service-group)# real-slb-server
```

```
10.1.1.3
Done!
P333R-LB-1(config-rsg:service-group)# exit
P333R-LB-1(configure)# virtual-server web-farm
Done!
P333R-LB-1(config-vs:web-farm)# id 1
Done!
P333R-LB-1(config-vs:web-farm)# vip 193.170.1.3
Done!
P333R-LB-1(config-vs:web-farm)# virtual-slb-service www-
service
Done!
P333R-LB-1(config-v-srvr:web-farm-v-slb-srvc:www)# id 1
Done!
P333R-LB-1(config-v-srvr:web-farm-v-slb-srvc:www)#
application tcp 80
Done!
P333R-LB-1(config-v-srvr:web-farm-v-slb-srvc:www)# rsg
service-group
Done!
P333R-LB-1(config-v-srvr:web-farm-v-slb-srvc:www)# exit
P333R-LB-1(config-vs:web-farm)# exit
```

Full NAT Load Balancing

In Full NAT based load balancing, the P333R-LB replaces the Virtual IP address of the P333R-LB with the real IP address of the Real Server, as in Half NAT load balancing. In addition, P333R-LB replaces the incoming client's IP address with a user-configured Proxy IP (PIP).

For the commands to configure Full NAT based load balancing, refer to Proxy IP (PIP) Commands on page 310.

Full NAT enables the P333R-LB to provide the following:

- **Eliminates the possibility of traffic that entered via the P333R-LB from bypassing the P333R-LB on return** - The PIP-client IP translation has to be performed by the P333R-LB. This ensures that even in complex networks, or when the Clients and Real Servers are in the same network, the return traffic from the Real Servers to the Clients always traverses the P333R-LB, and not through an potentially existing alternative route. This is also true in situations where the Real Servers are in the Public realm.

- **Single Virtual Service for multiple VIPs** - PIP enables mapping several VIPs to a single application. The Virtual Services, configured for each VIP, point to the same Real Server Group as well as to the same Real port. This enables emulating a configuration where different servers, each with a different IP address, provide the same service.
- **Direct Client-Real Server session access** - Full NAT enables sessions to the VIP as well as direct sessions between the Clients and Real Servers to traverse the same P333R-LB. The sessions to the VIP undergo Full NAT, while direct Client-Real Server sessions are routed by the P333R-LB without NAT. Such an application is not possible when implementing Half NAT.
- **Mapping different Virtual Services to the same port on the Real Servers** - Each service has its own PIP, enabling different flows to the same port.

You prepare banks of IP address ranges and associate each Virtual Service with a bank.

Note: You can create 64 banks of PIP addresses, with a total of 32,768 PIP addresses.

If you haven't configured a bank for a Virtual Service, P333R-LB uses Half NAT load balancing for that Virtual Service.

With Full NAT, responses made by the Real Server are forced to pass through the P333R-LB, where the Real Server IP is replaced with the Virtual Server IP, and the PIP with the client's original IP before forwarding them to the clients.

Note: Configure the PIP addresses on a dedicated subnet. If you configure a local interface on the PIP subnet, it is advertised throughout the network by the Routing Protocols as belonging to that router, like all other subnets. If no Local Interface is configured on the PIP subnet, Static Routes should be configured to reach that PIP subnet.

Note: The P333R-LB does not answer to ARP requests, or to pings on the subnet.

Note: You can not configure Full NAT for Active FTP Virtual Services.

Full NAT Based Configuration

Note: The following example uses sample IP addresses for illustration only.

Figure 14.9 illustrates a Full NAT based SLB configuration.

Figure 14.9 Full NAT Based SLB Configuration Example

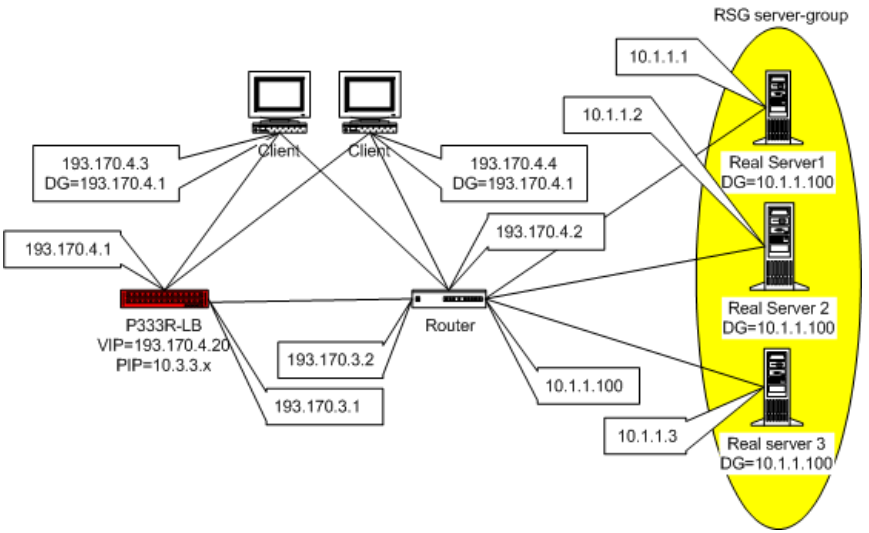

Figure 14.9 illustrates a sample configuration where two clients are connected to a P333R-LB which balances the traffic between three Real Servers. In addition to the traffic path through the P333R-LB, a direct path exists between the Clients and the Real Servers through another router. The P333R-LB is configured with a VIP of 193.170.4.20 and with a PIP bank on subnet 10.3.3.x.

In order to configure the load balancer according to Figure 14.9, perform the following commands:

```
P333R-LB-1(super)#interface router
Done!
P333R-LB-1(super-if:router)#ip address 193.170.3.1
255.255.255.0
Done!
P333R-LB-1(super-if:router)#exit
P333R-LB-1(super)#interface client
Done!
P333R-LB-1(super-if:client)#ip address 193.170.4.1
```

```
255.255.255.0
Done!
P333R-LB-1(super-if:client)#exit
P333R-LB-1(super)#ip default-gateway 193.170.3.2
Done!
P333R-LB-1(super)#real-slb-server 10.1.1.1
Done!
P333R-LB-1(super-r-slb-srvr:10.1.1.1)#exit
P333R-LB-1(super)#real-slb-server 10.1.1.2
Done!
P333R-LB-1(super-r-slb-srvr:10.1.1.2)#real-slb-server
10.1.1.3
Done!
P333R-LB-1(super-r-slb-srvr:10.1.1.3)#exit
P333R-LB-1(super)#rsg server-group
Done!
P333R-LB-1(super-rsg:server-group)#type-id slb 1
Done!
P333R-LB-1(super-rsg:server-group)#real-slb-server 10.1.1.1
Done!
P333R-LB-1(super-rsg:server-group)#real-slb-server 10.1.1.2
Done!
P333R-LB-1(super-rsg:server-group)#real-slb-server 10.1.1.3
Done!
P333R-LB-1(super-rsg:server-group)#exit
P333R-LB-1(super)#pip-bank 1
Done!
P333R-LB-1(super-pip-bank:1)#pip-addresses 10.3.3.2 10.3.4.2
Done!
P333R-LB-1(super-pip-bank:1)#exit
P333R-LB-1(super)#virtual-server slb
Done!
P333R-LB-1(super-v-srvr:slb)#id 1
Done!
P333R-LB-1(super-v-srvr:slb)#vip 193.170.4.20
Done!
```

```
P333R-LB-1(super-v-srvr:slb)#virtual-slb-service http
Done!
P333R-LB-1(super-v-slb-srvc:slb:http)#id 1
Done!
P333R-LB-1(super-v-slb-srvc:slb:http)#application tcp 80
Done!
P333R-LB-1(super-v-slb-srvc:slb:http)#pip-bank 1
Done!
P333R-LB-1(super-v-slb-srvc:slb:http)#rsg server-group
Done!
P333R-LB-1(super-v-slb-srvc:slb:http)#exit
P333R-LB-1(super-v-srvr:slb)#exit
P333R-LB-1(super)#
```

To insure that the packets destined to the PIP address reach the P333R-LB on the return path from the Real Server to the client, you must configure a Static Route in the router by entering: **`ip route 10.3.3.0 255.255.255.0 193.170.3.1 1 low`**

Note: The above configuration example stresses that even when there is a direct path from the router to the clients, with Full NAT the traffic is forced to traverse the P333R-LB for PIP-client IP translation. With Half NAT, in such a scenario, load-balanced sessions would have failed. This is because traffic from the Real Servers would have been routed directly to the clients, before reaching the P333R-LB. The clients expect a reply from the VIP, but instead would receive the reply from the real IP, and drop the packets.

Direct Server Return (DSR) (Triangulation) Redirection

In Direct Server Return (Triangulation) redirection, the redirection is based on replacing the destination MAC address with that of the Real Server. The Real Server is configured to send responses to the clients with the VIP as the source address. As a result, traffic is not required to traverse the load balancer.

Note: You can not implement DSR for Services using Full NAT, since the P333R-LB must replace the PIP with the original client IP for the returning packets.

DSR (Triangulation) Configuration Example

Note: The following example uses sample IP addresses for illustration only.

Figure 14.10 illustrates an SLB Triangulation configuration.

Figure 14.10 SLB Triangulation Configuration Example

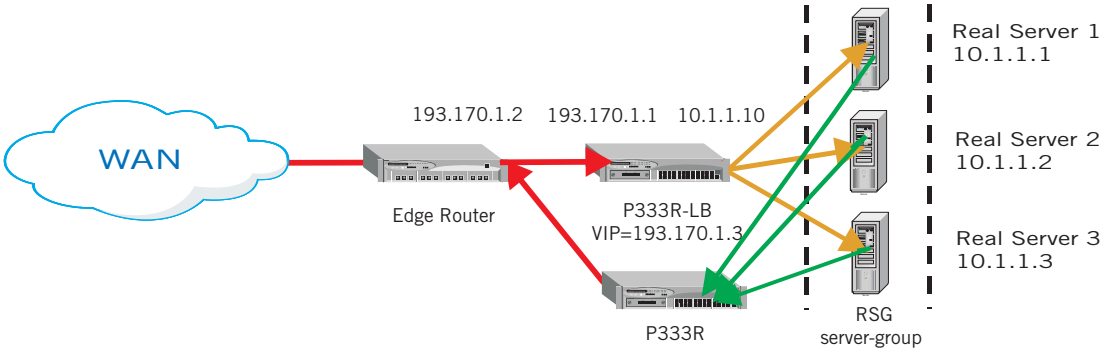

Note:
1. The VIP should be configured in the Real Servers as a "loopback" IP address.
2. Another router (like the P333R) may be configured as the default gateway of the Real Servers. This conserves resources and bandwidth on the P333R-LB that is tasked with balancing client requests.

The following configuration file is a result of configuring your network as shown in Figure 14.10:

```
hostname "P333R-LB"
!
interface "2"
 ip vlan name "Default"
 ip address 10.1.1.10       255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 193.170.1.1     255.255.255.0
!
ip default-gateway 193.170.1.2     1 low
```

```
!
real-slb-server 10.1.1.1
  direct-server-return
!
real-slb-server 10.1.1.2
 direct-server-return
!
real-slb-server 10.1.1.3
 direct-server-return
!
rsg "server-group"
 type-id slb 1
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
 real-slb-server 10.1.1.3
!
virtual-server "web-farm"
 id 1
 vip 193.170.1.3
 !
 virtual-slb-service "www-service"
  id 1
  application tcp 80
  rsg "server-group"
```

DNS Configuration

Note: The following example uses sample IP addresses for illustration only.

Figure 14.11 illustrates a DNS configuration, where DNS traffic to primary and secondary DNS servers is balanced.

Figure 14.11 DNS Configuration

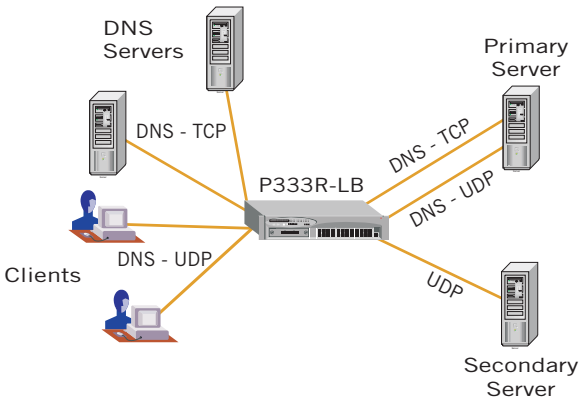

In Figure 14.11, DNS queries from clients arrive over UDP, while the DNS servers exchange DNS information via TCP. DNS includes a mechanism by which secondary and primary servers exchange information. The P333R-LB is required to balance UDP queries by clients across both primary and secondary servers, while forwarding TCP traffic only to the primary server. This is done by assigning two services to the Virtual Server that represents the DNS server to the world. The UDP service is mapped to a group of servers which include the primary and secondary Real DNS Servers. The TCP service is configured to include only the Real DNS server which has the primary role.

Combining Two Cache Applications on the Same Real Caches

Building on the fact that the P333R-LB allows for the same Real Servers to belong to different Load Balancing applications (i.e. Application Redirection and Server Load Balancing), it is possible to build a configuration that will implement transparent caching and non-transparent proxy caching on the same physical machines (refer to "Combining Two Cache Applications on the Same Real Caches" on page 178).

Load Balancing Metrics

There are different methods, or metrics, that a P333R-LB can use to distribute traffic among multiple Real Servers. A Real Server is a physical server that is associated

with a Real IP address. These metrics tell the router which server should receive each session.

SLB uses the following metrics:

- Round Robin.
- Hash.
- MinMiss Hash.

The default metric for SLB is Round Robin.

Each load balancing metric can be performed on one of the following metric parameters:

- SRC IP (source IP address).
- DST IP (destination IP address).
- SRC and DST IP (a combination of source and destination IP addresses).

In addition, each of the metrics may be weighted. For full information on weighted Real Servers, see Weighted Real Servers on page 40.

Round Robin

Round Robin is the default metric for SLB.

Using the Round Robin metric, new sessions are issued to each server in turn based on the Real Server weight. The first Real Server in the group receives the first 'n' sessions, where 'n' is the Real Server weight. The second Real Server receives the next 'n' sessions, and so on. When all the servers receive at least one session, the issuing process starts over with the first Real Server.

Hash

Using the Hash metric, sessions are distributed to Real Servers using a predefined mathematical hash function. This hash function is created by a hash value, performed on a specific parameter. The source and destination IP addresses are used as the hash function input.

P333R-LB creates a list of all the currently available servers. The result of the hash function is used to select a server from the list. Specific addresses always produce the same hash function, providing natural persistency.

If a server is added or removed to the group, the persistency will be broken. This occurs, since the order of the servers in the list changes, but the hash still points to the same list entries. The following figure illustrates how a loss of persistency occurs when a server becomes non-operational.

Figure 14.12 Hash Metric - Loss of Persistency

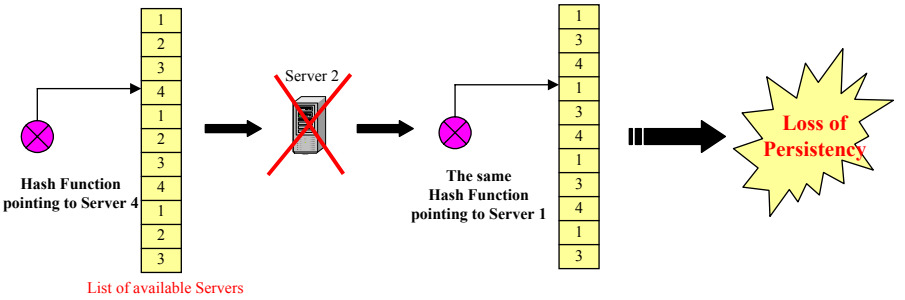

When Server 2 is removed from the group, the list of available servers is readjusted, causing a lack of persistency. However, if Server 2 becomes operational again, the list of available servers is restored to its original order, and persistency is recovered.

MinMiss Hash

MinMiss Hash distributes sessions to Real Servers in the same way as the Hash metric. However, MinMiss Hash sustains persistency even when a server is removed from the server group. When a server fails or is removed, the positions of the other Real Servers in the list do not change. Instead, the remaining servers are redistributed to the list entries freed by the failing server. The following figure illustrates how persistency is maintained even though a server becomes non-operational.

Figure 14.13 MinMiss Hash Metric - Persistency Sustained

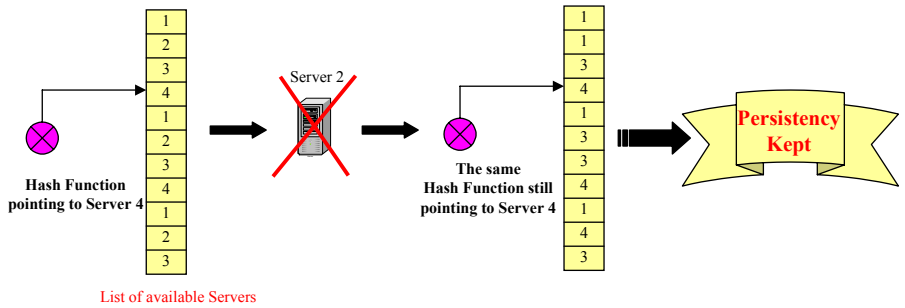

When Server 2 is removed from the group, the list of available servers is not readjusted. Only the list entries that are now empty are replaced with other available servers. Therefore, persistency is sustained for all available servers. If Server 2 becomes operational again, the list of available servers is recalculated so that the smallest number of servers is affected. However, the list of servers is not restored to its original configuration. As a result, persistency is only partially recovered.

Weighted Real Servers

You can assign weights to Real Servers to enable faster servers to receive a larger share of sessions. This minimizes overloading and maximizes functionality

If you assign a weight to a Real Server, the sessions are distributed to the servers in the same metric chosen (Hash or MinMiss Hash). However, weighted servers are assigned a larger share of sessions. For example, if you assign a weight of 100 to one and do not assign a weight to a second server, the weighted server receives 10 sessions for each session directed to the second server. This is because the default value of the server's weight is 10.

Health Check

The P333R-LB constantly health-checks the Real Servers to ensure that each Real Server (in this case, each server) is accessible and operational. A server that fails the health check is automatically removed from the load balancer's internal list of currently available servers.

P333R-LB supports the following health-check methods:

- **ICMP Echo** - Each server is periodically pinged and checked if an answer was received.
- **TCP Port Checking** - A TCP connection is periodically opened to every server, checking for successful completion of the connection.
- **HTTP Server Checking** - Useful for web applications, this method enables verifying HTTP server functioning by comparing the response from the server with a complete sample of the web page. The P333R-LB sends a request to the HTTP server to simulate an outside request. The P333R-LB compares the server's response with the sample to ascertain the level that the HTTP server is functioning. P333R-LB supports the following HTTP retrieving methods:
  - **Head** - Retrieve only the HTTP headers without document body content. This is the default method.
  - **Get** - Retrieve all data from the HTTP server.

  P333R-LB supports the following responses:
  - **Any response** - P333R-LB assumes the server is functional when P333R-LB receives any response.
  - **Any OK response** - P333R-LB assumes the server is functional when the server's response is one of the HTTP OK responses (codes 201-204).
  - **Exact OK response** - P333R-LB assumes the server is functional only when P333R-LB receives the exact response expected (code 200). This is the default method.

  You can configure the following parameters:
  - The port to access - if you do not configure a port, the default is in the following order: the Virtual Service Real port, Virtual port, or port 80.
  - HTTP Request type - GET or HEAD (default HEAD)
  - HTTP version 1.0 or 1.1 (default 1.0)

— Domain Name (for HTTP 1.1 only)

Note: If you selected HTTP 1.1, you can specify a Domain Name to be used in the Health Check query. If you don't specify a Domain Name, the Real Server's IP address and the Real port is used.

— URL to access - up to 255 characters (default "/")
— The expected response
— Expected response and type: Any Response, Any OK Response (HTTP status code 2xx), or Exact OK Response (HTTP status code 200) - default
— Expected String - a string up to 255 characters for comparison with the response
— Expected String Offset - an offset from the end of the TCP header (up to 1000 bytes) where the Expected String should appear

A successful HTTP Health Check is defined as one with both the expected response and the expected string valid.

- **Script Server Checking** - Supported over TCP, this method verifies the functioning of the server by running a script on the server. Script Health Check enables you to build your own script to run on the Real Server, and return a pre-defined response.
  You configure a complete and explicit request header as well as the expected response string with offset. The string's limiting factors are:
  — Maximum length of request header: 255 characters.
  — Maximum length of response string: 255 characters.
  — Maximum offset from end of TCP header: 1000 bytes.
  P333R-LB compares the server's reply with the expected reply you configured. You need to verify that the configured request results in the configured expected response.

Note: The default health check method for SLB is TCP Port Checking.

For the commands to configure the different Health Checks, refer to Health Check Commands on page 302.

Client Persistency

Persistency is a way to ensure that all traffic related to a given session and all sessions of a given characteristic are served by the same server.

Client persistency is the persistency between many sessions for one client. Client persistency ensures that all traffic from the client is directed to the same Real Server.

Client persistency is achieved either by using naturally persistent load balancing schemes (such as Hash or MinMiss Hash), or by forcing persistent load balancing decisions on non-persistent load balancing schemes (such as Round Robin). Decision forcing is performed by storing the history of the latest decisions in a cache for a limited time, and sending the packets to the appropriate server based on previous load balancing decisions.

Regardless of the client persistency nature of the selected load balancing metric, the P333R-LB offers a unique client persistency feature that is available in all load balancing metrics. Client persistency is based on a "persistency cache". Load balancing decisions are recorded in a persistency cache for a specified time configured by the user. When a new session that matches an entry in the persistency cache is processed by the P333R-LB, it is directed to the same server pointed by the cache (provided, of course, that the server is considered healthy).

The key to the persistency cache is based on the client IP, in combination with a wildcard. This allows persistency to be configured per an exact IP address, or per a group of addresses. For instance, in cases where clients hide behind a NAT device which selects NAT addresses from an address block of 255 addresses, enabling the persistency cache with a wildcard of 0.0.0.255 will map all clients to a single entry and a single Real Server.

Selecting a Load Balancing Metric and Persistency Option

The different load balancing metrics and persistency options provide you with flexibility.

Round Robin generally gives you the best load balancing solution.

Hash, with the key set to Src IP, gives a predictable metric, required when client persistency is needed and a redundant load balancer is used.

MinMiss Hash, with the key set to Src IP, gives the best persistency, without using the persistency cache.

The persistency cache mechanism enables you to specify the required client persistency time.

The wildcard feature allows the client to maintain persistency when the same client is mapped to a different IP address by a NAT device. However, persistency cache size is limited, and it somewhat increases the session processing overhead.

Port Re-mapping

The P333R-LB may be configured to re-map the destination port number when performing load balancing. For example, you might want to run the HTTP process on real servers using a different port number in order to enable more than one concurrent HTTP service.

Note:
1. In the P333R-LB, a Real Server can belong to multiple server groups as long as the groups are not running the same Virtual Service. If the groups are running the same service (e.g., HTTP), port re-mapping should be used.
2. Port Re-mapping should not be used in conjunction with Triangulation.

Application Redirection

This section provides information on Application Redirection (AR), and includes a configuration example.

With the importance of the Internet as an organization's source of information, normal operation of the LAN can be negatively impacted by congestion on the network router to the Internet.

Since much of the information retrieved from the Web is either repeatedly requested by a user or requested by multiple users, many organizations implement a local caching mechanism to prevent unnecessary WAN traffic.

The problem with local caches is that they must be on the traffic path between the client and the WAN router. As a result, all traffic, even non-cacheable traffic, passes through the caches.

The P333R-LB's Application Redirection (AR) redirects packets from their original destination to an alternative server, based on AR configuration. Since redirecting Web requests to local caches is the most common implementation of AR, it is also known as Cache Redirection.

The AR feature can also be used for policy-based (source-based) routing. For full details, see Policy-Based Routing (Source-Based Routing) on page 57.

Benefits

By redirecting client requests to a local cache or application server, you can increase the speed at which clients access the information and free up valuable network bandwidth.

AR for cache redirection provides the following benefits:

- Faster client access to information.
- Increased network bandwidth.
- Policy based routing.
- Only suitable traffic is directed to the local cache.
- Multiple caches can be connected and load-balanced.
- The redirection process is transparent to the client.
- Redundant caches can be configured.

How It Works

To enable AR, locate P333R-LB on the traffic route, instead of the local cache, and redirect packets from their original destination, the WAN access router, to alternative cache servers. The redirection process performs the following:

1 Checks whether the packet characteristics complies with one of the defined filter rules. You need to configure rules in order to define which clients/ destinations P333R-LB redirects to the cache applications.

2 Routes the packet to the cache server instead of the original destination, the

Web server.

3 The cache checks if it has the requested information. If it does, it replies to the client. If the cache does not have the information, it retrieves the information from the real Web server, and then replies to the client.

P333R-LB supports transparent caches. A transparent cache is capable of accepting packets not destined to its IP. The cache usually uses NAT in its IP stack, so the higher layers can process packets not destined to the cache IP.

The following steps illustrate an AR process:

1 The user issues an HTTP request with its IP address as the source address and the Web server's IP address as the destination address.

2 P333R-LB routes the packet to the Web cache(s) (load balancing if needed), but the packet still has the Web server's IP address as the destination IP address.

3 If the cache has the required page, the cache returns the page to the client with the destination IP address of the client and the source IP address of the Web server. If the cache does not have the required page, the cache returns the packet to P333R-LB, and it is routed to the Web server.

4 On returning from the Web server, P333R-LB routes the packet to the cache. The cache updates itself with the new page, and returns the packet to the client.

In the Cache Redirection packet flow, the following events need to be indistinguishable on the IP level:

- Redirection of packets from the client to the cache.
- Routing of packets from the cache to the WAN.

To address this issue, P333R-LB uses different VLAN areas for the clients and the cache. If a packet arrives to P333R-LB from the client's VLAN, it is redirected to the cache, and if it arrives from the server’s VLAN, it is routed to the WAN.

Configuration Examples

This section provides the following configuration examples:

- Application Redirection
- Combining Two Cache Applications on the Same Real Caches

Application Redirection

Figure 14.14 illustrates an AR configuration.

Figure 14.14 Cache Redirection Configuration Example

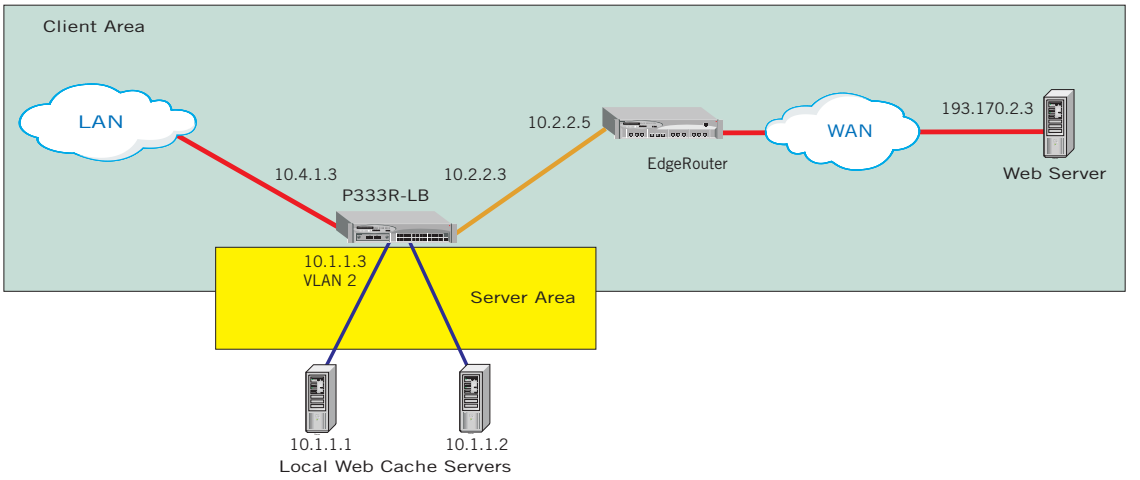

Note:

1. The Cache Servers must not be on the local subnet of one of the P333R-LB's local subnets.
2. The clients must not reside on the cache's subnet or VLAN.

In order to configure the load balancer according to Figure 14.14, perform the following commands:

```
P330-1(configure)# session router
Router-1(configure)# hostname "P333R-LB"
P333R-LB-1(configure)# set vlan 2 name v2
Done!
P333R-LB-1(configure)# interface 1
Done!
P333R-LB-1(config-if:1)# ip address 10.4.1.3 255.255.255.0
Done!
```

```
P333R-LB-1(config-if:1)# exit
P333R-LB-1(configure)# interface 2
Done!
P333R-LB-1(config-if:2)# ip address 10.2.2.3 255.255.255.0
Done!
P333R-LB-1(config-if:2)# exit
P333R-LB-1(configure)# interface 3
Done!
P333R-LB-1(config-if:3)# ip address 10.1.1.3 255.255.255.0
Done!
P333R-LB-1(config-if:3)# ip vlan 2
Done!
P333R-LB-1(config-if:3)# exit
P333R-LB-1(configure)# ip default-gateway 10.2.2.5
Done!
P333R-LB-1(configure)# set vlan-area 2 servers
Done!
P333R-LB-1(configure)# real-ar-server 10.1.1.1
Done!
P333R-LB-1(config-rsrvr:10.1.1.1)# exit
P333R-LB-1(configure)# real-ar-server 10.1.1.2
Done!
P333R-LB-1(config-rsrvr:10.1.1.2)# exit
P333R-LB-1(configure)# rsg cache-group
Done!
P333R-LB-1(config-rsg:cache-group)# type-id ar 1
Done!
P333R-LB-1(config-rsg:cache-group)# real-ar-server 10.1.1.1
Done!
P333R-LB-1(config-rsg:cache-group)# exit
P333R-LB-1(configure)# virtual-ar-service web-cache
Done!
P333R-LB-1(config-v-ar-srvc:web-cache)# id 1
Done!
P333R-LB-1(config-v-ar-srvc:web-cache)# application tcp 80
Done!
P333R-LB-1(config-v-ar-srvc:web-cache)# rsg cache-group
Done!
P333R-LB-1(config-v-ar-srvc:web-cache)# exit
P333R-LB-1(configure)# ar-filter 1 any any 1
Done!
```

The following configuration file is a result of the above configuration:

```
----------------
P333R-LB
----------------

set vlan       2 name "v2"
!#
!# Set the hostname of the device
!#
hostname "P333R-LB"
!
interface "3"
 ip vlan name "v2"
 ip address 10.1.1.3        255.255.255.0
!
interface "2"
 ip vlan name "Default"
 ip address 10.2.2.3        255.255.255.0
!
interface "1"
 ip vlan name "Default"
 ip address 10.4.1.3        255.255.255.0
!
ip default-gateway 10.2.2.5        1 low
!
real-ar-server 10.1.1.1
 !
real-ar-server 10.1.1.2
!
rsg "cache-group"
 type-id ar 1
 real-ar-server 10.1.1.1
 real-ar-server 10.1.1.2
!
virtual-ar-service "web-cache"
 id 1
 application tcp 80
 rsg "cache-group"
!
ar-filter 1 any any 1
```

```
!
set vlan-area 2 servers
```

Combining Two Cache Applications on the Same Real Caches

Based on the fact that P333R-LB allows for the same Real Servers to belong to different Load Balancing applications (i.e., Application Redirection and Server Load Balancing), it is possible to build a configuration that implements transparent caching and non-transparent proxy caching on the same physical machines, as shown in Figure 14.15.

In this case, it is required that the caches working as non-transparent proxy caches be configured as non-spoofing, i.e., caches that use their IP address as the source address. This is as opposed to "Spoofing" caches, which are capable of retaining the characteristics of the incoming packet even when they forward the packet further and keep the original Client IP address as the source address.

Figure 14.15 Two Cache Applications on the Same Real Caches

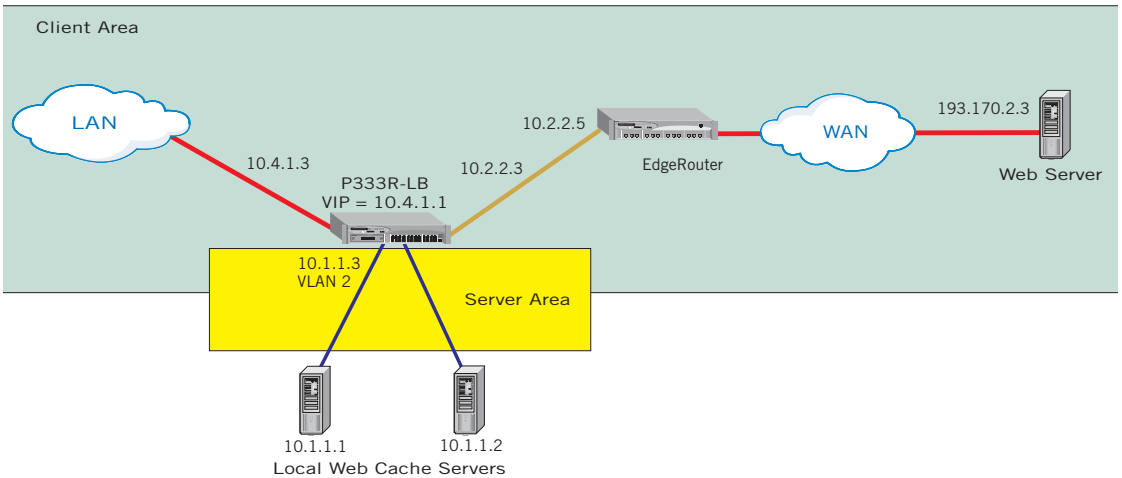

The following configuration file is a result of the above configuration:

```
set vlan       2 name "v2"
!
interface "3"
 ip vlan name "v2"
 ip address 10.1.1.3        255.255.255.0
!
interface "2"
 ip vlan name "Default"
 ip address 10.2.2.3        255.255.255.0
!
```

```
interface "1"
 ip vlan name "Default"
 ip address 10.4.1.3        255.255.255.0
!
ip default-gateway 10.2.2.5        1 low
!
real-slb-server 10.1.1.1
!
real-slb-server 10.1.1.2
!
real-ar-server 10.1.1.1
!
real-ar-server 10.1.1.2
!
rsg "none-transparent-proxy-group"

 type-id slb 1
 real-slb-server 10.1.1.1
 real-slb-server 10.1.1.2
!
rsg "transparent-proxy-group"
 type-id ar 1
 real-ar-server 10.1.1.1
 real-ar-server 10.1.1.2
!
virtual-server "none-transparent-proxy-server"
 id 1
 vip 10.4.1.1
 !
 virtual-slb-service "tcp-8080"
  id 1
  application tcp 8080
  rsg "none-transparent-proxy-group"
!
virtual-ar-service "transparent-proxy-server"
 id 1
 application tcp 80
 rsg "transparent-proxy-group"
!
ar-filter 10 any 10.1.1.0 0.0.0.255 no-ar
ar-filter 20 any any "transparent-proxy-server"
!
set vlan-area 2 servers
```

In Figure 14.15, the same two Cache Servers are configured as Real Servers for Server Load Balancing and for Application Redirection. Also, two Virtual Services are configured: one is a SLB service for the non-transparent proxy cache implementation, and the second is an AR service for the transparent cache implementation.

Traffic destined to the proxy cache, is sent by the client to the VIP as the Destination IP address, and dealt by the SLB Virtual Service (i.e., P333R-LB performs NAT on the packets and sends them to the Real Server based on the configured metrics). If the packets are to be forwarded to the Internet, P333R-LB receives the packets with the Source IP address of the cache (since it is a non-spoofing cache) and routes them to the Edge Router. On return, the packet is routed to the Real Server (since its IP address is now the Destination IP address) and the cache sends the packet back to the client.

Traffic not destined to the proxy cache, is sent with the Web Server's IP address (193.170.2.3) as the Destination IP address, and is processed by the AR Virtual Service as usual, based on the second rule of the "ar-filter" statement.

The reason for the first "ar-filter" statement in the above configuration (`ar-filter 10 any 10.1.1.0 0.0.0.255 no-ar`) is as follows: both in the SLB and in the AR cases it is possible that the packet is to be forwarded to the Web Server, if the required data is not in the cache. Upon returning, usually packets from the Web-server to the cache (in response to the non-transparent proxy cache SLB function) would be load-balanced according to the AR service metric. To prevent this from happening, the first "ar-filter" statement ensures that any packets destined to any of the Real Servers (caches) are not subjected to Application Redirection but rather are routed to the correct Real Server.

Load Balancing Metrics

There are different methods, or metrics, that a P333R-LB can use to distribute traffic among multiple cache servers. These metrics tell the router which cache server should receive the next session.

AR uses the following metrics:

- Round Robin.
- Hash.
- MinMiss Hash.

The default metric for AR is MinMiss Hash on the destination IP.

Each load balancing metric can be performed on one of the following metric parameters:

- SRC IP (source IP address).
- DST IP (destination IP address).
- SRC and DST IP (a combination of source and destination IP addresses).

In addition, each of the metrics may be weighted. For full information on weighted Real Servers, see Weighted Real Servers on page 54.

Round Robin

Using the Round Robin metric, new sessions are issued to each server in turn, based on the Real Server weight. The first Real Server in the group receives the first 'n' sessions, where 'n' is the Real Server weight. The second Real Server receives the next 'n' sessions, and so on. When all the servers receive at least one session, the issuing process starts over with the first Real Server.

Hash

Using the Hash metric, sessions are distributed to cache servers using a predefined mathematical hash function. This hash function is created by a hash value, performed on a specific parameter. The source and destination IP addresses are used as the hash function input.

P333R-LB creates a list of all the currently available cache servers. The result of the hash function is used to select a cache server from the list. Specific addresses always produce the same hash function, providing natural persistency.

If a cache server is added or removed to the group, the persistency is broken. This occurs, since the order of the cache servers in the list changes, but the hash still points to the same list entries. Figure 14.16 illustrates how a loss of persistency occurs when a cache server becomes non-operational.

Figure 14.16 Hash Metric - Loss of Persistency

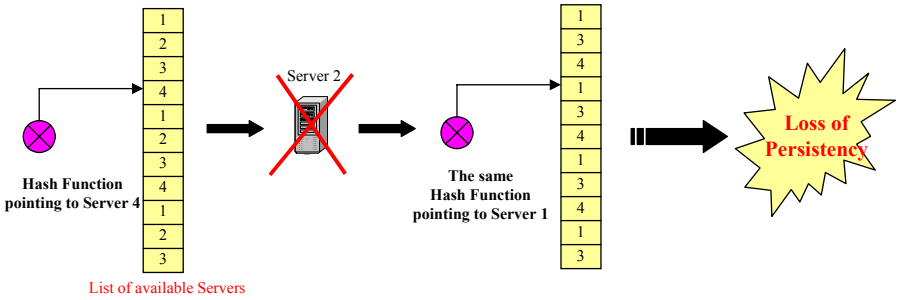

When Cache Server 2 in Figure 14.16 is removed from the group, the list of available cache servers is readjusted, causing a lack of persistency. However, if Cache Server 2 becomes operational again, the list of available cache servers is restored to its original order, and persistency is recovered.

MinMiss Hash

The default metric parameter for MinMiss Hash in AR is the Destination IP address.

MinMiss distributes sessions to cache servers in the same way as the Hash metric. However, MinMiss Hash sustains persistency even when a cache server is removed from the cache server group. When a cache server fails or is removed, the positions of the other cache servers in the list do not change. Instead, the remaining cache servers are redistributed to the list entries freed by the failing cache server. Figure 14.17 illustrates how persistency is maintained even though a cache server becomes non-operational.

Figure 14.17 MinMiss Hash Metric - Persistency Sustained

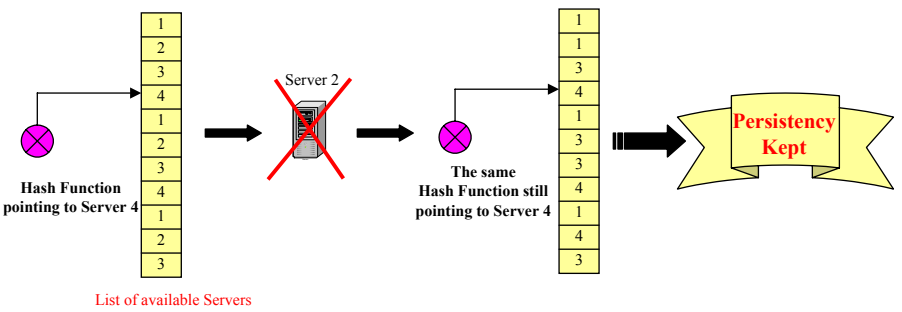

When Cache Server 2 in Figure 14.17 is removed from the group, the list of available cache servers is not readjusted. Only the list entries that are now empty are replaced with other available cache servers. Therefore, persistency is sustained for all available cache servers. If Cache Server 2 becomes operational again, the list of available cache servers is recalculated so that the smallest number of cache servers is

affected. However, the list of cache servers is not restored to its original configuration. As a result, persistency is only partially recovered.

Weighted Real Servers

You can assign weights to Real Servers (caches) to enable faster servers to receive a larger share of sessions. This minimizes overloading and maximizes functionality

If you assign a weight to a Real Server, the sessions are distributed to the servers in the same metric chosen (Hash or MinMiss Hash). However, weighted servers are assigned a larger share of sessions. For example, if you assign a weight of 100 to one and do not assign a weight to a second server, the weighted server receives 10 sessions for each session directed to the second server.

Health Check

P333R-LB constantly health-checks the Real Servers to ensure that each Real Server (in this case, each server) is accessible and operational. A server that fails the health check is automatically removed from the load balancer's internal list of currently available servers.

P333R-LB supports the following health-check methods:

- **ICMP Echo** - Each server is periodically pinged and checked if an answer was received.
- **TCP Port Checking** - A TCP connection is periodically opened to every server, checking for successful completion of the connection.
- **HTTP Server Checking** - Useful for web applications, this method enables verifying HTTP server functioning by comparing the response from the server with a complete sample of the web page. P333R-LB sends a request to the HTTP server to simulate an outside request. P333R-LB compares the server's response with the sample to ascertain the level that the HTTP server is functioning.
  P333R-LB supports the following HTTP retrieving methods:
  - **Head** - Retrieve only the HTTP headers without document body content. This is the default method.
  - **Get** - Retrieve all data from the HTTP server.

  P333R-LB supports the following responses:
  - **Any response** - P333R-LB assumes the server is functional when P333R-LB receives any response.
  - **Any OK response** - P333R-LB assumes the server is functional when the server's response is one of the HTTP OK responses (codes 201-204).
  - **Exact OK response** - P333R-LB assumes the server is functional only when P333R-LB receives the exact response expected (code 200). This is the default method.

You can configure the following parameters:

— The port to access - if you do not configure a port, the default is in the following order: the Virtual Service Real port, Virtual port, or port 80.
— HTTP Request type - GET or HEAD (default HEAD).
— HTTP version 1.0 or 1.1 (default 1.0).
— Domain Name (for HTTP 1.1 only).

Note: If you selected HTTP 1.1, you can specify a Domain Name to be used in the Health Check query. If you don't specify a Domain Name, the Real Server's IP address and the Real port is used.

— URL to access - up to 255 characters (default "/").
— The expected response.
— Expected response and type: Any Response, Any OK Response (HTTP status code 2xx), or Exact OK Response (HTTP status code 200) - default.
— Expected String - a string up to 255 characters for comparison with the response.
— Expected String Offset - An offset from the end of the TCP header (up to 1000 bytes) where the expected string should appear.
P333R-LB searches for the expected strings only in the first HTTP packet sent by the server as a response to the GET/HEAD request. If the string search fails, use the **`show hc last response`** command to view the received string and compare it with the expected string.

A successful HTTP Health Check is defined as one with both the expected response and expected string valid, as well as a successful completion of the TCP connection.

Note: In the HTTP Expected String Health Check script "\r\n" denotes "enter".

Refer to page 201 for a sample HTTP Health Check configuration.

- **Script Server Checking** - Supported over TCP, this method verifies the functioning of the server by running a script on the server. Script Health Check enables you to build your own script to run on the Real Server, and return a predefined response.
  You configure a complete and explicit request header as well as the expected response string with offset. The string's limiting factors are:
  — Maximum length of request header: 255 characters.
  — Maximum length of response string: 255 characters.
  — Maximum offset from end of TCP header: 1000 bytes.
  P333R-LB compares the server's reply with the expected reply you configured.

You need to verify that the configured request results in the configured expected response. P333R-LB searches for the expected string only in the first packet sent by the server as a response to the script query.
A successful Script Health Check is defined as one with a valid expected string, as well as a successful completion of the TCP connection.

Note: In the Script Health Check query and expected-string "\r\n" denotes "enter".

Refer to page 202 for a sample Script Health Check configuration.

Note: The default health check method for Application Redirection is Ping.

For the commands to configure the different Health Checks, refer to Health Check Commands on page 258.

Client Persistency

Persistency is a way to ensure that all traffic related to a given session, and all sessions of a given characteristic are served by the same server.

Client persistency is the persistency between many sessions for one client. Client persistency ensures that all traffic from the client is directed to the same Real Server.

Client persistency is achieved either by using naturally persistent load balancing schemes (such as Hash or MinMiss Hash) or by forcing persistent load balancing decisions on non-persistent load balancing schemes (such as Round Robin). Decision forcing is performed by storing the history of the latest decisions in a cache for a limited time, and sending the packets to the appropriate server based on previous load balancing decisions.

Regardless of the client persistency nature of the selected load balancing metric, P333R-LB offers a unique client persistency feature that is available in all load balancing metrics. Client persistency is based on a "persistency cache". Load balancing decisions are recorded in a persistency cache for a specified time configured by the user. When a new session that matches an entry in the persistency cache is processed by P333R-LB, it is directed to the same server pointed by the cache (provided, of course, that the server is considered healthy).

The key to the persistency cache is based on the client IP, in combination with a wildcard. This allows persistency to be configured per an exact IP address, or per a group of addresses. For instance, in cases where clients hide behind a NAT device which selects NAT addresses from an address block of 255 addresses, enabling the persistency cache with a wildcard of 0.0.0.255 maps all clients to a single entry and a single Real Server.

Policy-Based Routing (Source-Based Routing)

Using AR rules, you can configure your device to support Policy-Based Routing according to the client IP. In this case, you can select your next Hop according to the source IP configured on your device.

Selecting a Load Balancing Metric and Persistency Option

The different load balancing metrics and persistency options provide you with the following flexibility:

- Round Robin generally gives you the best load balancing solution.
- MinMiss Hash with the key set to Source IP gives the best persistency, without using the persistency cache.
- The persistency cache mechanism enables you to specify the required client persistency time.
- The wildcard feature allows the client to maintain persistency when the same client is mapped to a different IP address by a NAT device. However, persistency cache size is limited, and it somewhat increases the session processing overhead.

Chapter 15

Embedded Web Manager

This chapter describes the installation procedures for the Embedded Web Manager of the Avaya P330.

Overview

The Embedded Web Manager provides the following:

- Managing and monitoring Power over Ethernet.
- Device Configuration - Viewing and modifying the different device configurations.
- Virtual LANs - Viewing and editing Virtual LAN information.
- Link Aggregation Groups (LAGs) - Viewing and editing LAG information.
- Software Redundancy - Setting software redundancy for ports in an Avaya P330 Switch.
- Port Mirroring - Setting up port mirroring for ports in an Avaya P330 Switch.
- Trap Managers Configuration - Viewing and modifying the Trap Managers Table.
- Switch Connected Addresses - View devices connected to selected ports. Port Security.
- Intermodule Redundancy
  - — One pair per stack.
  - — Also operates as a result of a module fault, e.g., power failure.

System Requirements

Minimum hardware and Operating System requirements are:

- One of the following operating systems:
  - — Windows® 95
  - — Windows 98 SP1
  - — Windows 98 OSR (Second Edition)
  - — Windows ME
  - — Windows NT® 4 Workstation or Server
  - — Windows 2000 Professional or Server
- Pentium® II 400 Mhz-based computer with 256 Mb of RAM (512 Mb recommended)
- Minimum screen resolution of 1024 x 768 pixels
- Sun Microsystems Java™ plug-in version 1.3.1

- Microsoft® Internet Explorer® **or** Netscape Navigator/Communicator® (see table)

Table 15.1 Embedded Web Manager/Browser Compatability

	Windows 95 or NT	Windows 98, ME or 2000
Internet Explorer	5.0 or higher	5.01 or higher
Netscape Navigator/ Communicator	4.7	4.73

Note for users of Netscape Navigator: The Java plug-in requires certain services from **Windows 95** which are not present if **Internet Explorer** is not installed. In order to add these services to the operating system, please install Internet Explorer version 3 or higher. You can then use either browser to manage the switch.

Running the Embedded Web Manager

Note: You should assign an IP address to the switch before beginning this procedure.

1 Open your browser.
2 Enter the url of the switch in the format **`http://aaa.bbb.ccc.ddd`** where ***`aaa.bbb.ccc.ddd`*** is the IP address of the switch.

Note: The user name is “root”
The default password for read-write access is “root”.

Note: The Web management passwords are the same as those of the CLI. If you have created additional CLI user names or changed the default passwords then you can use those passwords for Web management as well.

The welcome page is displayed:

Figure 15.1 The Welcome Page

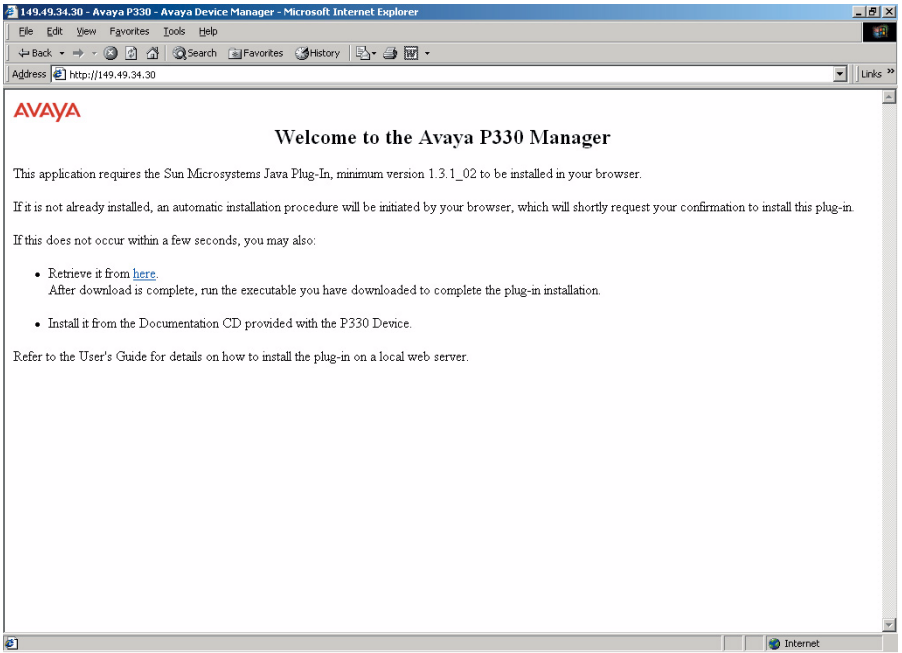
149.49.34.30 - Avaya P330 - Avaya Device Manager - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites History
Address http://149.49.34.30
Links
AVAYA
Welcome to the Avaya P330 Manager
This application requires the Sun Microsystems Java Plug-In, minimum version 1.3.1_02 to be installed in your browser.
If it is not already installed, an automatic installation procedure will be initiated by your browser, which will shortly request your confirmation to install this plug-in.
If this does not occur within a few seconds, you may also:
Retrieve it from here.
After download is complete, run the executable you have downloaded to complete the plug-in installation.
Install it from the Documentation CD provided with the P330 Device.
Refer to the User's Guide for details on how to install the plug-in on a local web server.
Internet

— If you have the Java plug-in installed, the Web-based manager should open in a new window (see Figure 15.2).

Figure 15.2 Web-based Manager

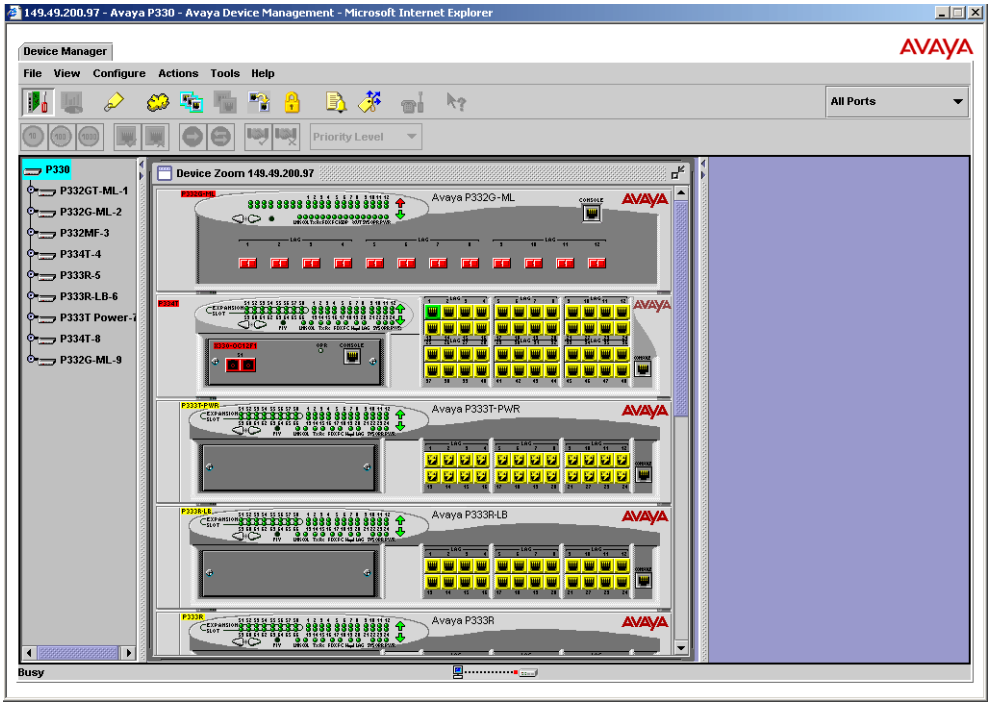

— If you do **not** have the Java plug-in installed, follow the instructions on the Welcome page that offers a variety of options to install the plug-in (see Figure 15.1).

Installing the Java Plug-in

If the network manager has configured the system, the plug-in should be installed automatically.

Note: Ensure that Java or JavaScript is enabled on your Web browser. Please refer to your browser on-line help or documentation for further information.

If the plug-in is not installed automatically, then you have three options for installing it manually:

Installing from the Avaya P330 Documentation and Utilities CD

1 Close all unnecessary applications on your PC.
2 Insert the "Avaya P330 Documentation and Utilities" CD into the CD drive.
3 Click **Start** on the task bar.
4 Select Run.
5 Type ***x:*****\emweb-aux-files\plug-in_1_3_1.exe** where ***x:*** is the CD drive letter.
6 Follow the instructions on screen.

Install from the Avaya Site

Click on the link in the Welcome page.

Install from your Local Web Site

Click on the link in the Welcome page.

Note: This option is only available if the network manager has placed the files on the local Web server.

Installing the On-Line Help and Java Plug-In on your Web Site

Note: This procedure is optional.

Copying the help files and Java plug-in to a local Web server allows users to access the on-line help for the Embedded Manager and enables automatic installation of the Java plug-in the first time the users tries to manage the device.

1 Copy the `emweb-aux-files` directory from the "Avaya P330 Documentation and Utilities" CD to your local Web server. Please refer to your Web server documentation for full instructions.

2 Define the URL in the Avaya P330 using the following CLI command:
`set web aux-files-url` ***`//IP address/directory name`***
where ***`//IP address/directory name`*** is the location of the directory from the previous step.

AVAYA P333R-LB

SECTION 2: TROUBLESHOOTING AND MAINTAINING THE P330

Troubleshooting the Installation

Troubleshooting the Installation

This section will allow you to perform basic troubleshooting of the installation. If you are unable to solve the problem after following the procedures in this chapter, please contact Avaya Technical Support. Refer to "How to Contact Us"for full details.

Table 16.1 Troubleshooting

Problem/Cause	Suggested Solution
Switch does not power up	
• AC power cord not inserted or faulty	• Check that the AC power cord is inserted correctly • Replace the power cord
If the cord is inserted correctly, check that the AC power source is working by connecting a different device in place of the P3330. • If that device works, refer to the next step. • If that device does not work, check the AC power	
• P3330 AC power supply not functioning	• Use an optional BUPS (Backup Power Supply) • Contact your local Avaya distributor. *The power supply is not user-replaceable.*
Stacking not functioning	
• X330-STK modules not inserted correctly (LEDs on stacking module do not light)	• Check that modules are installed correctly
• Octaplane™ cables not installed correctly (LEDs on stacking module do not light)	• Check that the cables are inserted correctly • Check that there are no cross-corrections
Expansion module not functioning	

Table 16.1 Troubleshooting

Problem/Cause	Suggested Solution
• Expansion module not inserted correctly	• Check that module are installed correctly

Maintenance

Introduction

This section provides basic maintenance information for the Avaya P330 switch and its components. For issues that are not covered in this chapter or in "Troubleshooting the Installation," please contact your Avaya representative.

Caution: Please refer to "Before You Install the P330" before undertaking any of the procedures detailed in this section.

Adding/Replacing an Expansion Sub-module

Caution: The expansion sub-modules contain components sensitive to electrostatic discharge. Do not touch the circuit board unless instructed to do so.

Adding an Expansion Sub-module to Avaya P330

1. Remove the blanking plate or other sub-module (if installed).
2. Insert the sub-module gently into the slot, ensuring that the Printed Circuit Board (PCB) is aligned with the guide rails.
   The PCB *not* the metal base plate fits into the guide rail.
3. Firmly press the sub-module until it is completely inserted into the Avaya P330.
4. Gently tighten the two screws on the front panel of the expansion sub-module by turning them.

Note: The Avaya P330 switch must not be operated with the expansion slot open; the expansion sub-module slot should be covered with the supplied blanking plate if necessary.

Replacing an Existing Expansion Sub-module

If an expansion sub-module is removed from the stack with the power supply on, all configuration definitions on expansion sub-modules are lost. Both procedures for replacing an expansion sub-module—with saving and without saving configuration definitions —follow:

Saving Configuration Definitions

1 Turn off the power supply.
2 Remove an expansion sub-module.
3 Insert another expansion sub-module.
4 Turn on the power supply.

Without Saving Configuration Definitions

1 Loosen the screws by turning the knobs.
2 Take hold of the two knobs (one near each side of the front panel) and pull gently but firmly towards yourself.
3 Insert another expansion sub-module or the blanking plate.

Replacing the Stacking Sub-module

To replace the X330STK stacking sub-module:

1 Power to the switch may remain on.
2 Loosen the screws to the stacking sub-module by turning the knobs.
3 Take hold of the two knobs (one near each side of the front panel) and pull gently but firmly towards yourself.
4 Insert the new stacking sub-module gently into the slot, ensuring that the metal base plate is aligned with the guide rails.
The metal plate—*not* the PCB of the X330STK— fits onto the guide rails.
5 Press the sub-module in firmly until it is completely inserted into the Avaya P330.

Caution: Ensure that the screws on the module are properly aligned with the holes in the chassis before tightening them.

6 Gently tighten the two screws on the side panel of the stacking module by turning the screws. **Do not use excessive force when tightening the screws.**

Updating the Software

This section provides the basic procedure for downloading and updating the P330 system software.

Caution: Please refer to "Before You Install the P330" before undertaking any of the procedures detailed in this section.

Software Download

You can perform software download using the CLI or Avaya UpdateMaster (part of the Avaya Multi-Service Network Manager Suite).

Obtain Software Online

You can download the firmware and Embedded Web Manager from the “Software Download” section at www.avaya.com/support.

Downloading Software

Download the firmware and Embedded Web Manager as follows:

Use the command in the Avaya P330 CLI:

`copy tftp SW_image <image-file> EW_archive <filename> <ip> <mod_num>`

image-file	firmware image file name (full path)
filename	Embedded Web Manager image file name (full path)
ip	The IP address of the TFTP server
mod_num	Target module number

Please see the CLI Chapters of the User’s Guides for related information.

Note: Please download both the new Avaya firmware and the new Embedded Web Manager versions. Whichever version of the firmware you decide to run, always be sure to match the correct firmware and Embedded Web Manager versions.

Download New Version without Overwriting Existing Version

Sometimes it is desirable to upgrade to a new software version while retaining the option of booting from the previous version. The following process copies the previous version from memory Bank B to Bank A, and download the new version to Bank B. This process accomplishes the following:

- prevents the embedded web image-file from being downloaded into Bank A - by providing a non-existant file name for the Embedded Web image file.
- preserves the old version in Bank A
- allows the user to boot from either Bank A or Bank B (i.e., using either the old or new software version)

Note: In normal operation, the Embedded Web file should be copied to Bank A, and the new software version should be downloaded to Bank B. This process copies the old software version to Bank A and the new software version to Bank B, and allows the user to boot from either version via the `set boot bank` command.

To perform this process:

copy tftp SW_image <new_ver_file> EW_image <dummy_file_name> <TFTP_server_IP_addr> <module_number>

Example:

copy tftp SW_image c:\versions\p330\p333t EW_image x 149.49.138.170 1

Note: Since file "x" doesn't exist the Embedded web image will not be downloaded.

How to Contact Us

To contact Avaya's technical support, please call:

In the United States

Dial 1-800-237-0016, press 0, then press 73300.In the EMEA (Europe, Middle East and Africa) Region

Country	Local Dial-In Number
Albania	+31 70 414 8001
Austria	+43 1 36 0277 1000
Azerbaijan	+31 70 414 8047
Bahrain	+800 610
Belgium	+32 2 626 8420
Belorussia	+31 70 414 8047
Bosnia Herzegovina	+31 70 414 8042
Bulgaria	+31 70 414 8004
Croatia	+31 70 414 8039
Cyprus	+31 70 414 8005
Czech Rep.	+31 70 414 8006
Denmark	+45 8233 2807
Egypt	+31 70 414 8008
Estonia	+372 6604736
Estonia	+372 6604736
Latvia	+371 721 4368

Country	Local Dial-In Number
Finland	+358 981 710 081
France	+33 1 4993 9009
Germany	+49 69 95307 680
Ghana	+31 70 414 8044
Gibraltar	+31 70 414 8013
Greece	+00800 3122 1288
Hungary	+06800 13839
Iceland	+0800 8125
Ireland	+353 160 58 479
Israel	+1 800 93 00 900
Italy	+39 02 7541 9636
Jordan	+31 70 414 8045
Kazakhstan	+31 70 414 8020
Kenya	+31 70 414 8049
Kuwait	+31 70 414 8052
Saudi Arabia	+31 70 414 8022

Country	Local Dial-In Number
Lebanon	+31 70 414 8053
Lithuania	+370 2 756 800
Luxemburg	+352 29 6969 5624
Macedonia	+31 70 414 8041
Malta	+31 70 414 8022
Mauritius	+31 70 414 8054
Morocco	+31 70 414 8055
Netherlands	+31 70 414 8023
Nigeria	+31 70 414 8056
Norway	+47 235 001 00
Oman	+31 70 414 8057
Pakistan	+31 70 414 8058
Poland	+0800 311 1273
Portugal	+351 21 318 0047
Qatar	+31 70 414 8059
Romania	+31 70 414 8027
Russia	+7 095 733 9055

Country	Local Dial-In Number
Slovakia	+31 70 414 8066
Slovenia	+31 70 414 8040
South Africa	+0800 995 059
Spain	+34 91 375 3023
Sweden	+46 851 992 080
Switzerland	+41 22 827 8741
Tanzania	+31 70 414 8060
Tunisia	+31 70 414 8069
Turkey	+800 4491 3919
UAE	+31 70 414 8036
Uganda	+31 70 414 8061
UK	+44 0207 5195000
Ukraine	+31 70 414 8035
Uzbekistan	+31 70 414 8046
Yemen	+31 70 414 8062
Yugoslavia	+31 70 414 8038
Zimbabwe	+31 70 414 8063

E-mail: csctechnical@avaya.com

In the AP (Asia Pacific) Region

Country	Local Dial-In Number
Australia	+1800 255 233
Hong Kong	+2506 5451

Country	Local Dial-In Number
Malaysia	+1800 880 227
New Zealand	+00 800 9828 9828

Country	Local Dial-In Number
Indonesia	+800 1 255 227
Japan	+0 120 766 227
Korea	+0 80 766 2580

Country	Local Dial-In Number
Philippines	+1800 1888 7798
Singapore	+1800 872 8717
Taiwan	+0 80 025 227

E-mail: sgcoe@avaya.com

In the CALA (Caribbean and Latin America) Region

E-mail: caladatasupp@avaya.com

Hot Line:+1 720 4449 998

Fax:+1 720 444 9103

For updated information, visit www.avaya.com/support and click “Global Support Organization (GSO)”.

© 2003 Avaya Inc. All rights reserved. All trademarks identified by the ® or TM are registered trademarks or trademarks, respectively, of Avaya Inc. All other trademarks are the property of their respective owners.